October 1993

Improved Inventory Models
for the United States Coast Guard Requirements Determination Process

CG201RD6

George L. Slyman
Dennis L. Zimmerman

Prepared for the U.S. Coast Guard pursuant to Department of Defense Contract MDA903-90-C-0006. The views expressed here are those of the Logistics Management Institute at the time of issue but not necessarily those of the U.S. Coast Guard or the Department of Defense. Permission to quote or reproduce any part except for Government purposes must be obtained from the Logistics Management Institute.

Logistics Management Institute
6400 Goldsboro Road
Bethesda, Maryland 20817-5886

Preface

Requirements determination is the process by which the Coast Guard supply system forecasts future customer demands and sets levels of inventory to satisfy those demands. Currently, the Coast Guard is in the process of modernizing its inventory-management system at its nonaviation supply centers. Modernization offers an opportunity to upgrade the current inventory models in the Coast Guard's requirements determination process by taking advantage of new computer technologies and modeling techniques that will help the centers continue to improve their performance. To support them, the Logistics Management Institute undertook a study to update the inventory models in the Coast Guard's requirements determination process.

Requirements determination spans a range of inventory-management decisions. Coast Guard item managers must decide how to forecast requirements, how to select methods to manage items, how to build inventory levels to satisfy demand, when to initiate repairs and procurement orders, and when to dispose of excesses. To assist them in making these decisions efficiently and effectively for thousands of items, they need a full set of inventory models.

Although different types of inventory-management models deal with different decisions, they are all interconnected. For example, the demand forecast from a forecasting model feeds into the computation of stockage levels performed by other models. If the objective is to improve requirements determination as a whole, the solution should address all of the models that are part of the process. It is not enough to improve demand forecasts if the levels computations that use those forecasts are not updated to capitalize on the improvements.

This report documents the LMI study of how the Coast Guard can update all of its inventory models. The report is intended to meet the needs of different readers:

- To those needing an overview, it informs and explains in easily understood terms the key components of an inventory management system.
- To those involved in inventory management, it describes and illustrates alternatives and extensions to the currently used inventory management models.
- To those analyzing performance and developing modernization requirements, it presents and discusses algorithms and decision logic for solving various inventory management problems.

Additionally, we make specific recommendations on which inventory models we believe the Coast Guard should use in its inventory-management system and we present those recommendations at the end of each chapter.

This report offers information in the following areas:

- *The requirements determination process.* It discusses the purpose and objective of the process, how inventory models contribute to the process, and the differences between requirements for consumable and reparable items. (Chapter 2.)

- *Forecasting.* We propose a demand forecasting methodology that includes 14 models, adaptations for nonrecurring and abnormal demand, and measures of forecast error. We also propose models for forecasting both procurement administrative lead time and production lead time and propose a method for forecasting repair cycle times. (Chapter 3; Appendix B describes the variety of stochastic models available for forecasting demand.)

- *Alternatives for managing items.* These alternatives range from local purchase to centrally-managed, not stocked to centrally-managed and stocked. We develop general cost equations for each alternative and offer a model for item management selection. (Chapter 4.)

- *Demand-supported stockage.* Specifically, we review computations for requisitioning objectives, economic order quantities, reorder points, procurement lead-time quantities, and safety-level quantities for all items as well as repair points and repair quantities for reparable items. (Chapter 5; readiness-based stockage, a potential alternative to demand-supported stockage, is discussed separately in Appendix A, and the algorithms available for determining an order quantity for a future demand pattern that is not constant are described in Appendix C.)

- *Insurance and low-demand stockage.* We offer a somewhat innovative response-based approach for managing low-demand items. (Chapter 6; Appendix D illustrates the approach using six low demand items.)

- *Inventory excesses.* We describe how to deal with excess inventory in the form of retention levels for supply center stocks and redistribution disposition models for unit excesses reported to supply centers. (Chapter 7.)

Contents

Contents (Continued)

Tables

Figures

Equations

Equations (Continued)

Equations (Continued)

CHAPTER 1

Introduction

To maintain the operations of its maritime fleet, the 5th largest in the world, and its far-flung aids-to-navigation stations, the Coast Guard depends on a network of Coast Guard and Department of Defense (DoD) supply centers, maintenance and supply depots, shipyards, and support activities; commercial manufacturers, vendors, and industrial sources; and vessel allowances. Its two nonaviation supply centers maintain an item inventory that was valued at almost $135 million on 30 June 1993. That inventory consists of the following items:

- Equipment and project materiel, e.g., diesel engines, vacuum pumps, field change kits, and initial outfitting sets
- Reparable components, subsystems, and assemblies, e.g., propeller shafts, combustion chambers, signal transmitters, and refrigerator compressor units
- Consumable repair parts, e.g., bolts, nuts, switches, heater elements, and fuses
- Bulk items and material, e.g., steel plate, cable, tubing, and pipe
- Expendable minor end items, e.g., life rafts and flags
- General-and special-purpose consumable items, e.g., ship's china, solar panels, and buoy-marking decals.

The supply centers maintain this inventory to satisfy the material requirements of its operating fleet of cutters, shore stations, and boat stations. The amount of inventory is based on forecasted customer demand, known material requirements (such as scheduled maintenance), and insurance/contingency stocks.

Currently, the Coast Guard is in the process of modernizing its inventory management system at its supply centers. In support of the modernization effort, the Logistics Management Institute (LMI) performed the inventory modeling study reported here. Specifically, the study looks at alternative methodologies the Coast Guard could use to forecast requirements, to select methods for managing items, to build inventory levels to satisfy demand, to initiate repairs and orders, and to dispose of excesses. It is aimed at modernizing the Coast Guard's two nonaviation supply centers – Supply Center Baltimore (SCB) and Supply Center Curtis Bay (SCCB) – although the methodologies could be of interest to the Coast Guard's Aircraft Repair and Supply Center.

Scope

Like transportation, distribution, procurement, storage, cataloging, and the other logistics disciplines, inventory management is well studied. Numerous management books and technical reports discuss supply, inventory control, materiel management, and the subtopic of requirements determination. Often those sources either provide general discussions of what is involved in determining requirements or present highly specialized decision models devoted to a particular aspect of the requirements determination process.

In this report, we do not offer any new theories that would contradict that body of knowledge. However, we do present in a single source a basic discussion of what is involved in the requirements determination process and what models are available to help determine requirements. We also make recommendations on what models the Coast Guard should consider for its supply management system.

Assumptions and Limitations

Three assumptions and one limitation govern the content of this report:

- "Change for the change's sake" is not part of the recommendations of this study. In fact, the Coast Guard supply centers are continually improving their level of support, and a legitimate question might be, "Why change the models that are used by the current requirements determination process if it is working?" Our answer is that the process can do better and it will need to in the future as more expensive high-technology items are introduced into the fleet and the fleet expects greater support while budgets grow at a slower pace. Modernization offers an opportunity to introduce state-of-the-art inventory models into the Coast Guard requirements determination process. If you are going to work with models, why not work with the best?
- In the past, the Coast Guard has not stressed inventory modeling in its supply management perhaps because of its engineering orientation. In the face of growing support and budget challenges, today's Coast Guard logisticians want to apply new computer technologies and modeling techniques to meet their goals. However, they are still interested in models that are understandable and that make sense.
- For the most part, future demands that Coast Guard units will place on a supply center are not known with any certainty. Demand itself is independent, i.e., a variety of units and processes generate demands for materials. Some supply center demand is like the demand in a manufacturing environment in which raw materials are needed to produce finished goods and the demand for these materials is dependent on the production schedule for the

finished goods. Other supply center demands are to meet unexpected needs or simply to be sure items are on hand when scheduled work comes due.

- In procuring material for stock or for customer use, Coast Guard buyers are governed by the Federal Acquisition Regulations (FAR). In considering commercial supply support alternatives for an item, the FAR sets limits on the contractual arrangements that can be made with vendor sources of supply.

CHAPTER 2

The Requirements Determination Process

PURPOSE

Requirements determination deals with satisfying customer orders for material. In the private sector, marketing people like to say that nothing really happens until a company gets a customer's order. Although that sounds good, it simply is not true in either the private or the public sector. In order to get the customer's order and to satisfactorily respond to it, many actions must occur before the order is received. Requirements determination is one of those actions.

As with any customers going to a store for supplies, personnel in Coast Guard units want to be able to go to their supply system and get what they need. In order for that to happen, the supply system must anticipate those needs and prepare accordingly. Requirements determination is the process by which the supply system forecasts future customer demands and sets levels of inventory to satisfy those demands when they occur.

Before we discuss how the Coast Guard could perform the requirements determination process, we should review what the Coast Guard wants to accomplish through the process. The two basic goals of any supplier are to provide maximum customer satisfaction and to maintain a minimum level of inventory. Maximum customer satisfaction occurs when all their potential orders are filled immediately. That level of satisfaction requires infinite or near-infinite levels of inventory. On the other hand, the minimum level of inventory is zero. Thus, the two objectives of a supplier are contradictory. Consequently, as shown in Figure 2-1, the requirements determination process must balance these two contradictory goals so that the system provides acceptable levels of customer satisfaction with affordable levels of inventory.

OBJECTIVE

In the private sector, the objective of any business is to maximize profits. That is how a businessperson measures success. Accordingly, that businessperson gears the requirements determination process to stock items that customers demand and that turn a profit when sold. However, profit is not the objective of the Coast Guard supplier; response-oriented fleet support is. Given that measure of success, the objective of the requirements determination process is to provide

a high level of supply support at the lowest cost. To do so may mean stocking slow-moving items as well as fast-moving items.

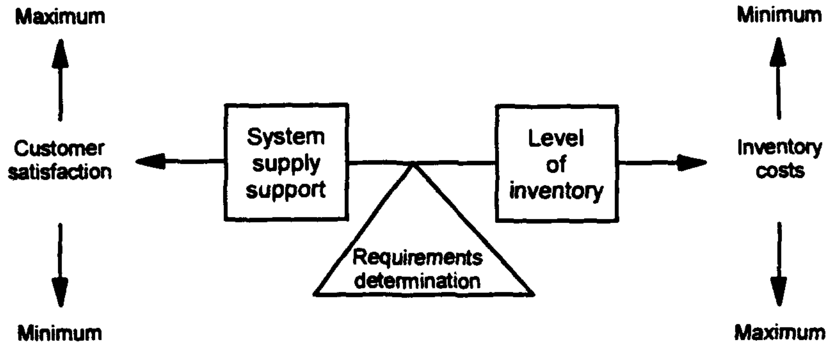

Figure 2-1.
Balancing Costs and Performance

Measures of Supply Performance

Since terminology is important to defining measures of supply performance, we introduce the terminology that we use through the very simplified diagram of supply support flows shown in Figure 2-2.

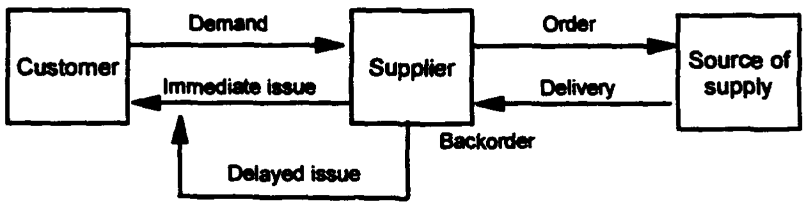

Figure 2-2.
Supply Support Flows

A customer's demand can either be filled immediately from on-hand stock or backordered when no stock is on hand and filled later with a delayed issue when replenishment stock is received.

The two principal measures of supply support are supply availability (or fill rate) and mean system response time (or average customer wait time). Supply availability is the percentage of customer demands that is filled immediately from on-hand stocks and is computed as follows:

$$Supply\ availability = \frac{Number\ of\ demands\ filled\ immediately}{Total\ number\ of\ demands} \qquad \text{[Eq. 2-1]}$$

Mean system response time is the average time needed to fill those demands that are filled immediately and those that are backordered and filled later with a delayed issue. It is computed as follows:

$$MSRT = k \bullet SA + B_t \bullet (1 - SA) \qquad \text{[Eq. 2-2]}$$

where

$MSRT$ = mean system response time

k = average time to fill a demand immediately (usually assumed to be zero)

SA = supply availability

B_t = average time to fill a backordered demand (i.e., a demand not filled immediately)

It is worth noting that supply availability is a component of mean system response time.

An Example

Consider a simple example of a supplier who is managing two items. Item 1 has 90 demands and Item 2 has 10. The supplier is able to fill Item 1 demands 100 percent of the time from on-hand stock. The supplier does not stock Item 2 but has a source for filling backorders in 15 days. The supply availability in this example is 90 percent [90/(90+10) = 0.90]. The mean system response time would be 1.5 days [0 • 0.90 + 15 • (1-0.90) = 1.5].

If the supplier increased costs by stocking Item 2 to a level at which 50 percent of its demand were filled immediately, supply availability would increase to 95 percent and mean system response time would decrease to 0.75 day. On the other hand, if the supplier decided to continue to not stock Item 2 but found an alternative source of supply that would fill backorders in 7.5 days, the mean system response time would again decrease to 0.75 day although supply availability would remain at 90 percent. If the supplier decreased costs by reducing the stock for Item 1 to a level at which 90 percent of its demand was filled immediately and the other backordered 10 percent was filled from a source of supply in 20 days, supply availability would decrease to 81 percent [(90 • 0.9)/100 = .81]. The average time to fill a backorder would be

17.4 days [(20 • 9 + 15 • 10)/(9+10) = 17.4] and the mean system response time would increase to 3.3 days [0 • 0.81 + 17.4 • (1-0.81) = 3.3].

Responsiveness — the Coast Guard Measure of Performance

By definition, mean system response time is a more comprehensive measure of supply support than supply availability since it considers how well the system responds to all demands. Time on backorder is important to the customer whose demand is on backorder. For that reason, the Coast Guard has adopted response time as a measure of supply support for the future and injects the responsiveness theme into its integrated logistics strategic planning.

However, unlike supply availability, mean system response time has no intuitive goals (i.e., where a supply availability goal of 95 percent was better than 85 percent and represented an "A" effort, is an "A" effort in terms of response time 3, 7, or 20 days?). Moreover, response time measure is somewhat counterintuitive in that supply support is improved by decreasing the goal not increasing it (i.e., a goal of 5 days is better than a goal of 10 days). As the Coast Guard suppliers perform their future requirements determination processes, they will need to work with operating units to establish desirable and affordable response time goals.

Operational Availability as a Measure of Performance

In closing our discussion of objectives for the requirements determination process, we need to touch on a new measure of supply performance that is primarily emerging from work being done within the DoD. Often referred to as readiness-based sparing, that work entails developing inventory models that have as their objective the operational availability (A_o) of weapon systems. These models are much more complex than supply-availability or response-time models and require considerably more data. (Appendix A describes a general readiness-based sparing model.) Moreover, the computation of A_o involves not only supportability factors but also reliability factors (e.g., mean time between failures) and maintainability factors (e.g., organizational maintenance repair times).

To illustrate A_o with our previous example, we will assume that the two items (Item 1 and Item 2) constitute an equipment system, only one equipment system exists, a single echelon of supply supports the equipment system, system downtime is solely a function of item backorders and every backorder downs the equipment system, and demands are for one of each item. We can compute the expected backorders (EBOs) for an item at any point in time as the number of backorders multiplied by the time on backorder (B_t) divided by 365 days in a year. If we compute the equipment system's A_o as the product of the probabilities that each item is causing the system to be down (i.e., 1 minus the EBOs for an item), then we have the data shown in Table 2-1.

Table 2-1.
Comparing Supply Performance Measures

Case	Item settings				Results from previous supply performance analysis		Results of A_o analysis		
	Item 1 (90)		Item 2 (10)		Supply availability	Response time (days)	EBO		A_o equal to (1-EB01) x (1-EB02)
	Avail.	B_t	Avail.	B_t			Item 1	Item 2	
1	100%	-	0%	15	90%	1.5	0	0.41	59%
2	100%	-	50%	15	95%	0.75	0	0.21	79%
3	100%	-	0%	7.5	90%	0.75	0	0.21	79%
4	90%	20	0%	15	81%	3.3	0.49	0.41	30%

Table 2-1 demonstrates the results of different performance computations and also the strong relationship between response time and A_o. In general, response time and A_o move in the same direction since the two are related. That is not always true of supply availability; situations can be constructed in which it increases while A_o decreases.

The A_o approach is particularly appealing because it attempts to directly relate supply to customer operations, especially to those of equipment system managers. However, thus far, its actual implementation has been limited. As noted before, readiness-based sparing requires significantly more data than conventional sparing. Moreover, the algorithms are much more complex and require a very strong mathematical background to understand. Given these added requirements, the question arises whether it is worthwhile to adopt an A_o approach when a response time approach may mimic the A_o results.

Currently, the Coast Guard Aircraft Repair and Supply Center is moving toward an A_o-oriented requirements process (notably, aviation support is where A_o models have had their greatest success in DoD). At some future time, if SCB and SCCB were to consider changing from response time to an A_o as the measure of supply support, each should perform a cost and benefit analysis to justify the added costs of the readiness-based sparing approach.

The Role of Models

To provide responsive support to the fleet, Coast Guard supply centers must address the inventory problems of how to forecast demand and lead times, how to determine the best method for managing an item, and how to set inventory levels for stocked items. If supply centers managed 1 to 10 items, the solutions to these problems could well revolve around qualitative procedures that are labor

intensive. However, since they manage thousands of items, they must rely on inventory models in their requirements determination process.

Inventory models consist of mathematical equations that provide a logical framework for solving inventory problems. Typically, such models define the solution variables associated with inventory decisions in terms of problem variables and their interrelationships. When values are assigned to the problem variables, the models will yield their solutions.

For example, if we decide that the quarterly demand forecast for an item should be the average of its demand over the past 2 years, the associated 8-quarter moving average model would be given by:

$$D_f = \frac{\left(\sum_{i=1}^{8} D_i\right)}{8} \qquad \text{[Eq. 2-3]}$$

where

D_f = the forecasted demand

D_i = quarterly demand observations

Σ = summation over index i

i = index from 1 to 8;

that is, the forecasted demand equals the sum of the last 8 quarterly demand observations divided by 8. If the observed demands for the last quarters were as follows

Quarter	1	2	3	4	5	6	7	8
Demand	20	22	640	0	25	17	15	21

then the sum of the demand over the 8 quarters would be 760 and the forecasted demand would be 95.

As our example illustrates, models are ideally suited for computers, which are designed to perform mathematically computations quickly. Their speed combined with their capability to consider any number of variables and consistently produce dependable answers makes computers highly desirable in calculating inventory requirements.

However, models are only tools and they have limitations. An inventory manager looking at our demand history may know that Quarters 3 and 4 are one-time aberrations and should not be used when predicting demand. Using

the remaining 6 quarters of data, the inventory manager arrives at a more correct forecast of 20.

CONSUMABLE-VERSUS-REPARABLE REQUIREMENTS

Consumables are commonly defined as items that are thrown away after they are used. In contrast, reparables are items that are repaired for reuse after they are used. In terms of the supply support diagram, the difference between consumables and reparables is that reparable items have two sources of supply instead of one. Specifically, the source of supply for consumable items is procurement while the sources of supply for reparable items are procurement and maintenance (see Figure 2-3).

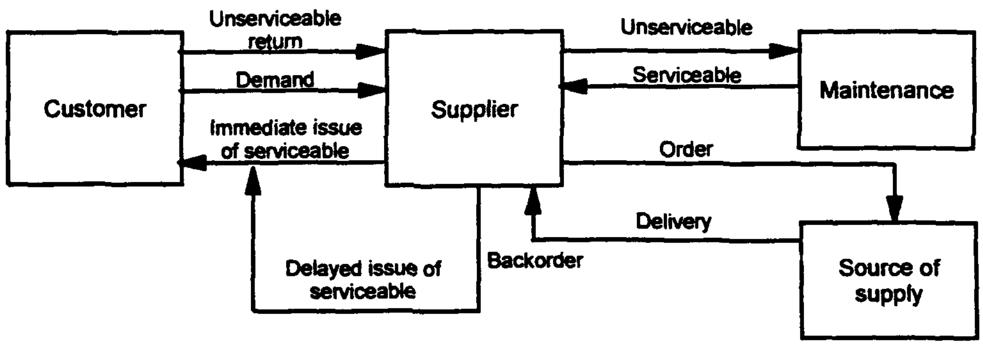

Figure 2-3.
Reparable Supply Support Flows

The determination of which items are reparables and which are consumables can either be based on an engineering judgement or on a level-of-repair analysis (LORA). The former is generally used for low-cost items where repair/no-repair experience exists with similar items. The latter is used for high-cost items where the question is not only whether or not to repair the item but also where to repair it. A LORA develops a total cost equation for each alternative representing a combination of repair/no-repair decisions for different levels of maintenance and selects the low cost solution.

In terms of the requirements determination process, more inventory management decisions are associated with reparable items than consumable items. In particular, decisions of when, how much, and how long to repair affect reparable inventory levels. (One approach to scheduling repair is to direct it whenever an unserviceable is turned in; another is to compute repair quantities and time levels and use them like order quantities and reorder points.)

CONCLUSIONS AND *RECOMMENDATIONS*

- Requirements determination is the supply center process for planning and providing for customer orders from the fleet in a manner that produces a high level of responsive support at the lowest possible cost.

- Because both procurement and maintenance are sources of supply for reparable items, the requirements determination process for reparable items is more multifaceted than that for consumable items, whose sole source of supply is procurement.

- Inventory models consist of mathematical equations that provide a logical framework for solving the inventory problems related to the requirements determination process, namely, forecasting customer needs and supplier resupply times, determining how to manage an item, setting levels of stocks, and deciding when to retain and/or redistribute stock.

- The Coast Guard has adopted a responsiveness objective in its integrated logistics strategic planning; as a measure of supply support, mean system response time considers both the number of times an order is not filled immediately and the time required to fill a backorder.

We recommend that SCB and SCCB use mean system response time as their measure of support and the objective of their inventory models, for two reasons:

- For equipment-related items, unlike supply availability, it tracks with equipment operational availability and avoids the extensive data and technical requirements of models that have equipment operational availability as their objective.

- For items that are not equipment related, it addresses the customer's question "how long is it going to take to satisfy my need?"

CHAPTER 3

Forecasting

GENERAL DISCUSSION

Requirements determination starts with forecasting a customer's needs and a supplier's resupply times. Those forecasts are the basis for determining whether to manage an item as stocked or nonstocked and, for stocked items, how much to stock. However, no universal model exists for forecasting demand or resupply times. Rather, a wide variety of models exist for use in different forecasting situations. This chapter discusses these models and recommends those that the Coast Guard should use to forecast demand, procurement lead time, and repair times.

Classification of Forecasting Methods

In their book, *Statistical Methods for Forecasting,* Abraham and Ledolter classify forecasting methods. We describe their classification in Table 3-1.

Table 3-1.
Methods of Forecasting

Method	Description
Qualitative (or subjective)	A nonrigorous, usually not reproducible approach that relies on intuitive, largely educated guesses that may or may not depend on past data.
Quantitative	A rigorous, reproducible approach that relies on mathematical or statistical models.
Deterministic models	The relationship between the variable being predicted and the variable(s) used to make the prediction is exact and known with certainty (e.g., a traditional "law" of physical science – $E = MC^2$).
Probabilistic or stochastic models	The relationship between the variable being predicted and the variable(s) used to make the prediction is not exact and is not known with certainty but is inferred from past data (e.g., the typical relationship between the demand for a product and its price – the higher the price, the lower the demand).

To forecast requirements for a large number of items, quantitative methods are needed. However, if an item manager (IM) has information that could improve the accuracy of a forecast, the forecasting system must have the

capability to supplement or override the quantitative forecast with that information.

Quantitative Forecasting Models

In general, quantitative forecasting models produce forecasts that are based on either future program requirements or on historical data for an item. Program-based forecasting models can either be deterministic or stochastic while forecasting models that are based on historical data are stochastic.

An example of a deterministic program-based forecast taken from the initial outfitting of Coast Guard recruits would be the demand forecast for duffle bags. If the Coast Guard plans for 200 recruits in the coming month and each recruit receives one bag, the forecasted demand for duffle bags would be 200 bags for the coming month.

An example of a stochastic program-based forecast would be failure rate for engine components. If an engine valve had a history of failing every 1,000 hours (a statistical average) and the engine program for the coming year is 100,000 hours, the failure forecast for the item would be 100 failures for the coming year. Accuracy in such a case depends on the accuracy of the usage rate (i.e., the mean time between failures) and the program data (i.e., the projected hours of use in the engine program).

A variety of stochastic models use historical data. One is the moving average model currently used by the SCB and SCCB to forecast demand. Another very popular model for forecasting demand in inventory systems is the single exponential smoothing model.

The historical data that go into stochastic models are usually ordered in some time sequence referred to as a time series (e.g., the four past quarters of demand for an item, which are collected in data buckets, constitute a time series). The pattern of the time series serves as the basis for selecting the forecasting model that will produce the minimum error. Appendix B provides a brief description of the following categories of stochastic models as well as general descriptions of individual models in the category and how they are formulated:

- Simple time series models
- Decomposition method
- Smoothing models
- Seasonal smoothing (Winter's) model
- Linear trend models

- Box-Jenkins method
- Nonlinear trend models
- Combined forecasts.

Given the variety of models that exist, the question is which model or models should the Coast Guard select for each of its various forecasting needs. In their book, *A Managerial Guide to Business Forecasting,* Ellis and Nathan list criteria for selecting a forecasting model. We describe their criteria in Table 3-2.

Demand Forecasting

Since demand forecasts are at the center of most stockage decisions, their accuracy is key to the success of a supply system. If forecasts are higher than actual usage, the result is excess inventory; if forecasts are lower than actual usage, the result is excessive backorders and an increased possibility of not meeting response-time goals.

The automated systems at the Coast Guard supply centers forecast future demand on the basis of historical demand. However, the IM can ultimately override an automated forecast with a manual forecast based on available information on future demand.

As previously noted, the Coast Guard, like many other Federal supply activities, uses a moving average technique to forecast demand. Other commonly used forecasting models are those that involve smoothing where an item's forecast is updated periodically with a weighted average of the old forecast and the demand for the last period. Variations revolve around the value of the weight (or smoothing constant) used for the last period's demand, the period between forecasts, the use of a filter to screen out unusually high demands from the forecast, and the extent to which nonrecurring demand is included in the forecast.

Since a new item has no historical demand, its forecast is developed from a manufacturer's or engineering estimate. An item's initial forecast is phased out as demand develops, normally over the first few years of its life in the system.

Application

Typically, the requirements determination process uses the demand forecast to decide whether the item should or should not be stocked and if stocked, what its inventory level should be. Our analysis of the Coast Guard's inventory management requirements and processes yielded the following matches of these

Table 3-2.
Criteria for Selecting a Forecasting Model

Criteria	Description
Pattern of data	The basic patterns of data are horizontal (constant/level pattern), trend (upward or downward sloping pattern), seasonal (short-term upward and downward repeating patterns), and cyclical (long-term upward and downward repeating patterns). *Pattern* — *Models/methods* horizontal — simple models, smoothing models, and Box-Jenkins method trend — linear/nonlinear models, Winter's model, and Box-Jenkins method seasonal — simple model, Winter's model, decomposition, and Box-Jenkins method cyclical — decomposition and Box-Jenkins method
Forecast horizon	The forecast horizon can be immediate, short-term, medium-term, or long-term — only the last three are involved in requirements determination. For short term forecasting (1 to 3 months), the trend of the time series is generally not important. For medium-term forecasting (3 months to 2 years), the trend, seasonal, and cyclical aspects of the time-series pattern are all important. For long-term forecasting (over 2 years), trend is very important.
Use of forecast	The use of the forecast governs the time available for forecasting and the understanding required for acceptance. For requirements determination, time limitations rule out the more complicated modeling methods such as decomposition and Box-Jenkins.
Forecast cost	When selecting a model, the cost of operating it and meeting its data requirements need to be considered. Simple and smoothing models have the smallest total cost while trend models are between the high-cost Box-Jenkins method and the low-cost single exponential smoothing model.
Single product or product lines	If a single item is being forecasted, the more complicated, more time-consuming techniques may be appropriate. However, where forecasts are needed for many items instead of a single product, the concern for computer time and ease of updating promote the use of either the simple models or the smoothing models.
Forecast accuracy	Measures of forecast error, such as mean square error (MSE) or mean absolute deviation (MAD), can be used to select the model that minimizes the forecast error. Combining forecasts from different models tends to reduce error.

stockage decisions with the Ellis and Nathan criteria for selecting a forecasting model (see Table 3-2):

- *Data* - Coast Guard demand is independent[1] and item demand patterns have level/horizontal components, trends, seasonal variations, cyclical movements, and random fluctuations.
- *Horizon* - Generally the stockage decisions have medium-term horizons (i.e., 6 months to 2 years) except for retention stocks, which have a long-term horizon (i.e., greater than 2 years).
- *Use* - The forecasts are automated or at least the initial recommendation to the IM is automated.
- *Cost* - The forecasts must be made at a low cost without the input of highly specialized expertise from professional forecasters.
- *Single or product line* - The forecasts must be made for a number of items.
- *Accuracy* - Error is always involved in forecasting independent demand.

Proposed Methodology

Given the application of the demand forecast, we believe that the Coast Guard does not need the more complicated and expensive forecasting techniques of the decomposition and Box-Jenkins methods. However, in light of the growing availability of computer power, we also believe that the Coast Guard can move away from the use of a single model to the use of multiple models. We believe that the use of multiple models will reduce forecast error by providing models that are more closely tailored to the demand patterns of individual items.

Figure 3-1 illustrates our proposed demand forecasting methodology.

Forecasting with a set of models and selecting the model whose result has the greatest accuracy is known as focus forecasting. While that approach is not a new idea, the growth in computer power has made it feasible, whereas in the past excessive time was required to compute a number of forecasts using different models for a large number of item.

---

[1] In inventory modeling, the distinction between dependent and independent demand is the difference between demand that can be quantified by (i.e., depends on) scheduled events such as those found in a manufacturing or repair process and demand that may be related to scheduled events but cannot be quantified by them.

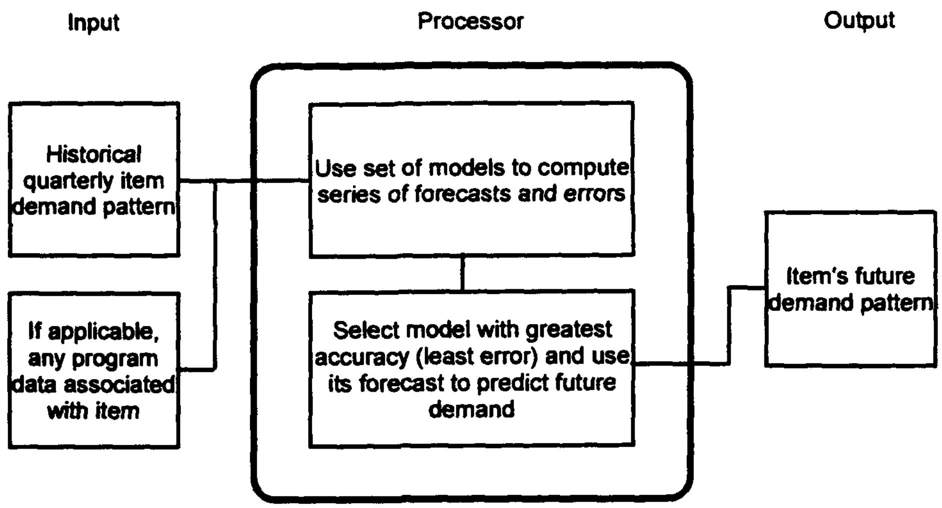

Figure 3-1.
Proposed Demand Forecasting Methodology

A somewhat new idea is the preparation and storage of a future demand pattern instead of a single future demand forecast. To prepare the pattern, the selected model first produces a forecast for the next pattern (the first period in the pattern). The model then uses that forecast as the next observation in building another forecast for the second period in the pattern. The process continues in that way until the last period in the pattern is forecasted.

That approach has the following advantages over the often used approach of straight-lining the single forecast into the future:

- It allows the forecaster to include historical trend, seasonal, and cyclical factors in forecasting future demand.
- It allows the IM to input known future demand information directly into the forecast.
- It produces the same forecast as straight-lining for items with generally level patterns.

These advantages are illustrated in Figure 3-2.

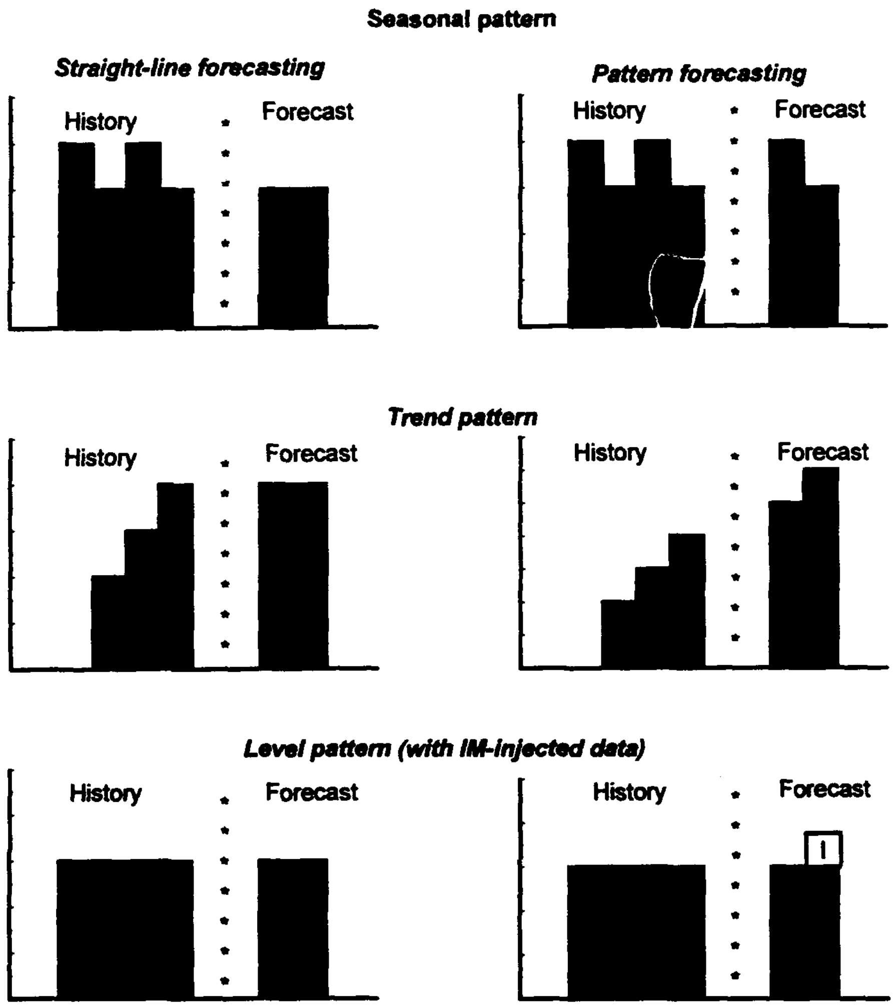

Figure 3-2.
Straight-line Versus Pattern Forecasting

TREATMENT OF NONRECURRING DEMAND

Coast Guard units can code a demand as nonrecurring (rather than recurring) if they feel that it is a one-time requirement. However, as a supply center receives these demands over time from different units, a recurring pattern may emerge. Since the inclusion or exclusion of nonrecurring demand in demand forecasts will raise or lower most inventory levels, respectively, their treatment in forecasting is important.

Figure 3-3 illustrates demand patterns in which the inclusion of nonrecurring demand can either increase or decrease the "lumpiness," of the pattern and consequently either increase or decrease the forecast error.

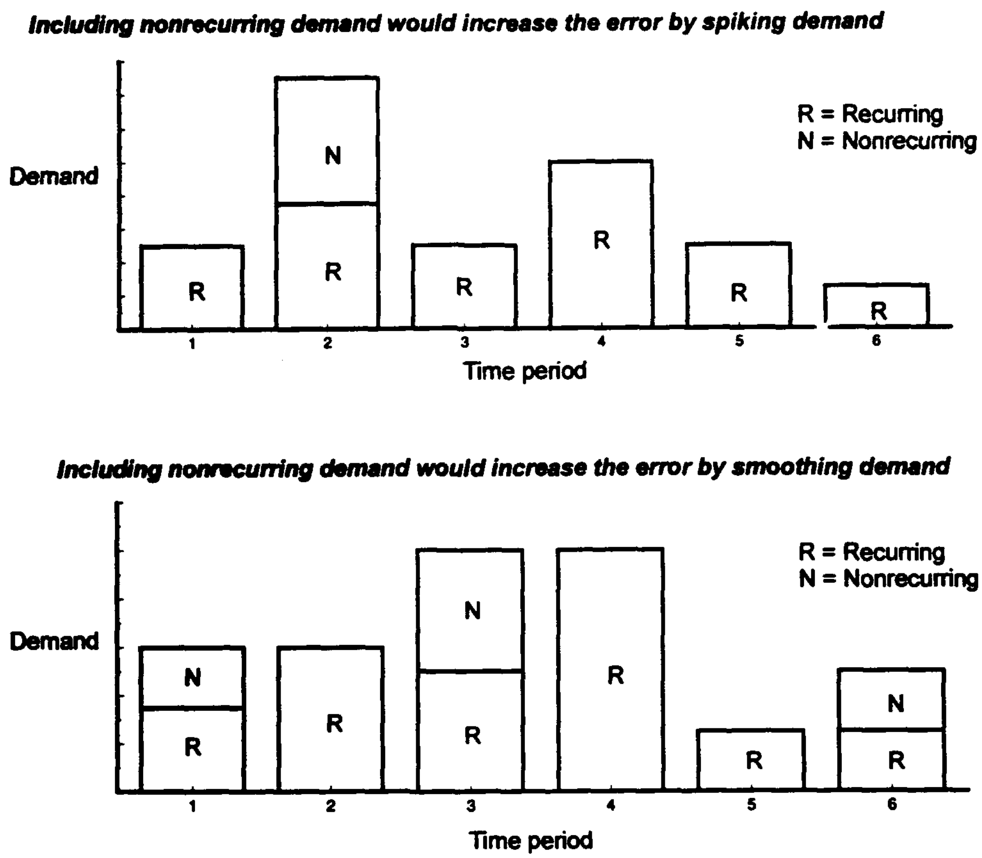

Figure 3-3.
Increasing or Decreasing Forecast Error with Nonrecurring Demand

Five basic alternatives exist for handling nonrecurring demand. The advantages and disadvantages of each are shown in Table 3-3.

We believe a combination of the last two alternatives best suits the Coast Guard. That is, for low-priced items, include all demand when forecasting and for high priced items, compute dual forecasts and select the best one with or without nonrecurring demand. That approach saves computer time and avoids any costly mistake in building a false level for a high-priced item.

DAMPENING

Dampening or filtering the demand that goes into a forecast has proven to be an effective technique in reducing forecast errors. As noted in the previous discussion on nonrecurring demand, unusual spikes in demand cause forecasting models to overforecast and that results in inflated or false inventory levels. Dampening consists of computing the normal dispersion of the data and using that dispersion to identify data that lie outside the norm, i.e., that are "outliers."

Table 3-3.
Alternatives for Treating Nonrecurring Demand

Alternative	Advantages	Disadvantages
Exclude all nonrecurring demand	For some items, avoids spikes caused by one-time data	Fails to account for all demand in building levels, and for some items, misses smoothing caused by repeated nonrecurring demand
Include all nonrecurring demand	Accounts for all demand in building levels and, for some items, increases smoothness of demand caused by repeated nonrecurring demand	For some items, builds false levels because of one-time spikes in demand
Include a percentage of nonrecurring demand	Accounts for most demand in building levels	Fails to take full advantage of smoothed patterns for some items and fail to fully avoid spikes for other items
Include all for items with low prices and exclude those with high prices	Accounts for all demand in building levels for low-cost items and avoids potential false levels for high-cost items	For some high-cost items, fails to account for repeated nonrecurring demand
Compute forecast with and without nonrecurring demand and use the one with least error	Improves forecast accuracy and builds levels on a pattern that is most representative of an item's future demand	Requires additional computer time to run dual sets of forecasts and select the best model

Dampening then reduces any outlier to the norm so that it is not in the historical demand pattern that goes into a forecasting model.

One approach to dampening is to compute the standard deviation (SD) of the data and use a multiple of it as a norm for identifying and reducing any outlier. To illustrate how this dampening works, we put together the example shown in Table 3-4 for a $2,909 life raft.

We believe that dampening generally reduces the forecast error, but it must be used judiciously so that it does not cloak real, dramatic changes in demand. We suggest that the Coast Guard use a dampening factor of two standard deviations for high-priced items (i.e., conservative forecasting when the cost of an error could be high) and three standard deviations for low-priced items (i.e., less conservative forecasting when the cost of an error would be less). Dampening should only be a temporary change to the historical demand data for an item for purposes of forecasting with a time series model. We caution against dampening the historical demand that goes into computing a usage rate for the usage model if those data are correlated with equipment operating data. However, if they are

Table 3-4.
An Example of Dampening Demand

Demand Pattern	Quarters								Comment
	1	2	3	4	5	6	7	8	
Original, undampened series	16	37	3	3	7	0	3	0	Mean = 8.625 and SD = 11.085
Dampened series, 3 standard deviations	16	37	3	3	7	0	3	0	No change since all values are between 0 and 41.88
Dampened series, 2 standard deviations	16	31	3	3	7	0	3	0	Quarter 2 reduced to mean plus 2 SDs
Dampened series, 1 standard deviation	16	20	3	3	7	0	3	0	Quarter 2 reduced to mean plus 1 SD

not correlated, i.e., they move in opposite directions in a period of high demand or demand goes up but operations stay the same, dampening is appropriate.

Recommended Models

In our proposed demand forecasting methodology (see Figure 3-1), we recommended a process known as focus forecasting, i.e., using a set of models and selecting the one that results in the least error. We believe that as a minimum, the set of demand forecasting models should include the 14 models and combinations of models shown in Table 3-5. Appendix B gives examples of how each of the time-series models (i.e., models 1 through 7 and 9 through 12) computes a forecast.

Overall, models in this proposed set are simpler and are slanted towards smoothing or averaging demand. The simpler models were chosen because numerous studies have shown them to be just as effective (or more effective in some cases) as the more complex models in forecasting demand. The slant toward smoothing evolves from the fact that most demand patterns tend to be random. The large number of models is an attempt to take advantage of what each model brings to forecasting and to take advantage of combined forecasts, which often tend to outperform single ones.

If an IM cannot associate an item with an equipment system and its operational program, only 11 models rather than 14 would be run for that item. For example, if an item was solely used on a Coast Guard cutter, had an average of 30 demands per quarter for the past two years, and this year's cutter program remains the same as last year's except for the last quarter during which steaming

Table 3-5.
Recommended Set of Demand Forecasting Models

Number	Model	Rationale
1	The basic model (BAS)	The simplest model; one that works well with items that experience strings of zero demand quarters.
2	A basic 4-quarter seasonal model (SBAS) (a year-ago model)	The simplest seasonal model.
3	A 4-quarter moving average model (MA4Q)	A simple model that tends to smooth down random fluctuations in demand.
4	An 8-quarter moving average model (MA8Q)	The longer period increases the smoothing of random demand
5	A single-exponential smoothing model with alpha = 0.1 (SES1)	Another simple model that tends to smooth demand but gives more weight to the most recent demand.
6	A single-exponential smoothing model with alpha = 0.2 (SES2)	The larger alpha increases the weight placed on the most recent demand.
7	A linear regression model (REGR)	A simple trend model.
8	A simple usage model for items associated with a program (USAGE)	Adds to the set a model that considers the future program when forecasting demand.
9	BAS and MA8Q	Combines smoothing and last demand.
10	SBAS and MA8Q	Combines smoothing and seasonality.
11	SBAS and SES2	Combines smoothing and seasonality.
12	SBAS, MA8Q, and REGR	Combines smoothing, seasonality, and trend.
13	MA8Q and USAGE	Combines smoothing and future program.
14	SES1 and USAGE	Combines smoothing and future program.

Note: The alpha in models 5 and 6 is the smoothing constant.

hours are to be cut in half, a usage model would project demands of 30 for the next 3 quarters and 15 for the fourth quarter.

Although we feel that the recommended set of models is an excellent starting point for forecasting demand, it is not meant to be an exhaustive list that should never be changed. We believe the Coast Guard should establish a test "laboratory," in which it could use actual demand patterns from a representative set of items to evaluate the recommended set and any additional models that would improve the set. Improved accuracy, understandability, and ease of use should guide the selection of additional techniques.

Input Requirements

The data requirements for the recommended set of models are shown in Table 3-6. The Coast Guard already collects demand data in quarterly "buckets,"

so that requirement is not new. It also retains the most recent forecast and the past quarter's demand so that requirement is not new. However, the storage of program data in the supply system is new.

Table 3-6.
Data Requirements for Forecasting Demand

Requirement	Models
Buckets to store past demand (8 buckets required to accommodate all models)	Basic model, basic seasonal model, moving average models, and combinations of these models
Most recent forecast and most recent demand bucket	Single exponential smoothing models and combinations with them
Past and future program data	Usage model and combinations with it

Program data can take on several forms: they can, for example, be information on the steaming-hour program for items associated with ships or, perhaps, on the number of overhauls for items associated with the maintenance of a principal end item (e.g., a cutter or a buoy). The simple usage model would combine past program data with past demand to compute a usage rate (e.g., 4 overhauls in a year with 15 units demanded would translate to a usage rate of 3.75 units per overhaul). Then, it would combine future program data with the usage rate to compute a forecast (e.g., a projected program of 5 overhauls in the coming year would produce a forecast for each quarter of 4.69 units, assuming the overhauls are evenly distributed over the year).

As presented, the USAGE model is a simple linear model, which is adequate when missions and programs remain fairly stable. However, when dramatic changes occur with regard to missions, such as cutting the length of a typical cruise in half, past usage rates may not be representative. When that occurs, the IM must either rule out the model or obtain another source for the usage rate.

Processing Forecasting Requirements

The three elements of processing forecasts are resources, frequency, and range of items. For the simple time-series models we recommend, no special forecasting expertise is needed. Consequently, the only resource needed to do the forecasting is computer time. That time will depend on the speed of the computer and the number of items processed.

Forecast Frequency

A common misconception is that in forecasting, like in many other activities, more is better. Since demand forecasts predict random events, they will always

be wrong. The variety of customers and circumstances generating demand in the Coast Guard cause demand to have a degree of randomness, which, in turn, causes forecasts to have some degree of error. Consequently, if we increase the forecasting frequency in an attempt to better measure how demand is moving and to use that information to improve our future demand forecasts, we will actually increase the forecast error that comes from randomness.

On the other hand, if forecasts are not updated on a regular basis, they will lag trends and this will also increase the forecast error. In short, we have the spectrum shown in Figure 3-4.

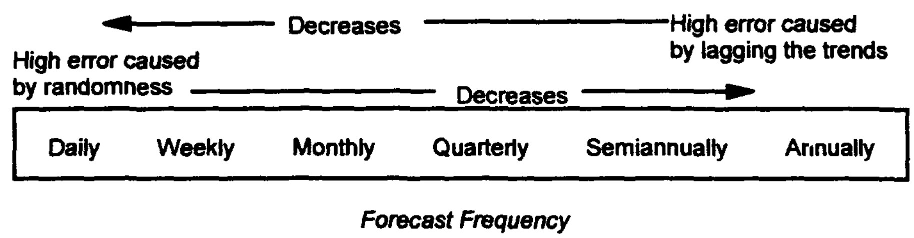

Figure 3-4.
Forecast Frequency and Error

In view of this apparent dilemma, we believe that the Coast Guard should adopt a middle-of-the-road stance, that is, forecast on a quarterly basis to reduce the influence of randomness and to avoid lagging the trends.

Item Screening

There is no question that a demand forecast should be prepared for an item that experiences regular demand; the question is how to test an item experiencing little or no demand. Computing a forecast for an item with no demand would be a waste of time since the forecast would be zero. If an item had only one quarter of demand, no matter if it was one requisition for a quantity of one or 10 requisitions for a total quantity of 50, any forecast derived from this sparse time series would be questionable.

However, not every item will experience demand every quarter. In fact, our review of Coast Guard items found that many had one or more quarters of zero demand. So the question is, when do future demand or future demand patterns generated by forecasting models become of such little value that a new approach is needed?

We believe that forecasting models should be used if demand for an item has been recorded for two or more quarters in the past eight quarters of demand. The forecast for an item with no demand should be zero; an item with only one quarter of demand should be considered a low-demand item and managed

under rules not requiring an individual item demand forecast. (We discuss such items in Chapter 6.)

Selecting a Model Based on Forecast Error

Focus forecasting, or multiple model forecasting, hinges on the ability to select the best forecasting model for an item. Ideally, we want the model that will give us the least error in the forecast. Just as we use historical demand to predict future demand, we are going to use historical error to predict future error. The problem is that several statistics exist for measuring forecast error.

At this point, we want to distinguish between bias and accuracy in forecasting. If we look at the differences between the actual and forecast demands and sum those differences over time, we are measuring the bias of the model, that is, we are determining whether it is overforecasting (negative sum) or underforecasting (positive sum). If we look at the size of the differences without regard to whether they are negative or positive, we are measuring accuracy. The norm in selecting a demand forecasting model has usually been forecast accuracy.[2] We support that approach to judging forecast error because we believe that in inventory management, an overforecast is just as undesirable as an underforecast.

Measures of Forecast Error

Using actual demand for a period (*Di*) and forecasted demand for the same period (*Fi*), we can compare differences over some number of periods (*n*) in terms of the following statistics:

- Mean error (ME)

$$ME = \frac{\sum_{i=1}^{n} (D_i - F_i)}{n} \qquad \text{[Eq. 3-1]}$$

- Mean percentage error (MPE)

$$MPE = \frac{100}{n} \sum_{i=1}^{n} \frac{(D_i - F_i)}{D_i} \qquad \text{[Eq. 3-2]}$$

- Mean square error (MSE)

$$MSE = \frac{\sum_{i=1}^{n} \left(D_i - F_i\right)^2}{n} \qquad \text{[Eq. 3-3]}$$

[2] A recent large-scale forecasting study conducted by the Defense Logistics Agency also considered bias (see Bibliography for reference). However, the recommended models from that study are similar to those we recommend.

- Mean absolute error or mean absolute deviation (MAD)

$$MAD = \frac{\sum_{i=1}^{n} |D_i - F_i|}{n} \qquad \text{[Eq. 3-4]}$$

where the vertical lines on both sides of the difference means the absolute (unsigned) value of the difference.

- Mean absolute percentage error (MAPE)

$$MAPE = \frac{100}{n} \sum_{i=1}^{n} \left| \frac{D_i - F_i}{D_i} \right| \qquad \text{[Eq. 3-5]}$$

The first two statistics (ME and MPE) measure bias while the others measure accuracy.

Of the three accuracy statistics, we believe the MSE is the most suitable. Although all three measure the unsigned size of the error, and by squaring the error, the MSE places more emphasis on large forecast errors. Consider the case in which two forecasting models have the same sum of absolute errors but one has consistent differences while the other has both smaller and larger differences. We would want the one with the consistent differences because for any single period we have a lower risk of making a large mistake. For that reason, we prefer the MSE over the MAD or MAPE as a measure of forecast accuracy.

Utilizing Forecast Error to Select a Model

The four-step procedure we propose for using the forecast error to select a model for an item is as follows:

1. Using the item's procurement lead time,[3] determine the number of quarters to use to compute the error as follows:

Lead time	Number of quarters
≤ 6 months	2
> 6 and ≤ 9 months	3
> 9 months	4

2. For each of the 14 models (11 if the USAGE model does not apply), use the number of quarters of demand data determined in Step (1) and the same number of forecasts for that item (each forecast for a model must be stored when it is made) to compute an MSE for the model.

[3] We are building our error forecast on the lead time because, as we discuss in Chapter 5, the error over the lead time is the basis for setting items' safety stocks.

3. Find the lowest MSE and select that model.

4. Optionally, if all the signs (i.e., + or −) of the 2 to 4 differences between demands and forecasts are the same for the model with the lowest MSE, find the second lowest MSE and, if it is close to the lowest MSE (say within 10 percent) and not all of its signs are the same, select that model.

To illustrate this procedure, we constructed a simplified example that only considers the first 6 models instead of all 14 (or 11) models. If the item in our example has an 8-month lead time, three quarters would be used to compute the error (see Table 3-7). In our example, the four-quarter moving average (Model 3) has the lowest MSE (119.0) and the signs of the differences are not all in one direction. Therefore, we would select it.

Table 3-7.
Simplified Example of Forecast Model Selection

Model	Past qtr.	Demand	Forecast	Difference	Squared	Sum	MSE
1. BAS	1	12	10	2	4		
	2	40	12	28	784		
	3	23	40	−17	289	1,077	359.0
2. SBAS	1	12	50	−38	1,444		
	2	40	20	20	400		
	3	23	10	13	169	2,013	671.0
3. MA4Q	1	12	20	−8	64		
	2	40	23	17	289		
	3	23	21	2	4	357	119.0
4. MA8Q	1	12	19	−7	49		
	2	40	19	21	441		
	3	23	19	4	16	506	168.7
5. SES1	1	12	15	−3	9		
	2	40	15	25	625		
	3	23	17	6	36	670	223.3
6. SES2	1	12	17	−5	25		
	2	40	16	24	576		
	3	23	21	2	4	605	201.7

Output Requirements

We previously stated our preference for pattern forecasting instead of straight-line forecasting. Applying that preference requires that a number of future demand buckets be created to store forecasts over some time horizon. Just

as data buckets are required to store historical demand data, another set of buckets would store future demand predictions.

Since requirements determination decisions mostly involve medium time horizons, a 2-year, or 8-quarter, future time horizon should be adequate. How the predictions that go into these buckets are generated depends on which of the 14 potential models is selected as the best model for a particular item. In what follows, we discuss how the future forecasts are generated by each of the models.

Generation of Future Demand Forecasts

If the selected model is either the basic model or one of the single exponential smoothing models, the 8 future quarterly demand buckets are filled with the same number, namely, the forecast from the model because the weighted or smoothed average of the same number yields that number. Therefore, the future demand pattern generated by a single exponential smoothing model is level. And, any demand projections past the 2-year period are determined by straight-lining.

If the selected model is the basic seasonal model, one of the moving average models, or the linear regression model, then the eight buckets are filled by the iterative method previously described. The future demand pattern is not level. Any demand projections past the 2-year period have to be handled cautiously as the continuation of a statistically derived seasonality or trend may be questionable. As a conservative alternative, we suggest straight-lining the average of the last 4 quarters.

We can illustrate the results of the iterative process using our previous example where a 4-quarter moving average model was selected as the forecast model. For that example, the future demand pattern would be as follows:

Future Quarter:	1	2	3	4	5	6	7	8	>8
Forecast:	21	24	27	24	24	25	25	25	25

In this case, our straight-lining suggestion for quarters after the eighth quarter coincides with what happens if the moving average is continuously applied in the future.

If the selected model is the simple usage model, then the eight buckets are filled by applying the usage rate and program for each quarter. Depending on the program data, the future demand pattern could be level or not. Any demand projections past the 2-year period are generated by the same usage-rate/program-data procedure. Additional quarterly buckets are required to span the period of time covered by future program data.

If the selected model is a combination not previously mentioned, eight forecasts for each model in the combination are first computed using the previously

described procedures. Then they are averaged and put into their respective future demand buckets. For projections past the 2-year period, straight-lining is used as previously described except in the case of combinations involving the USAGE model. For those combinations, an average of the straight-line forecast and usage-rate/program-data forecast is used along with additional quarterly buckets.

Inclusion of Known Future Customer Demand

Occasionally, the fleet provides an IM information pertinent to future demand for an item. For example, the IM may learn of a major overhaul of cutters that will require 120 units of an item above its normal demand in the second quarter of the coming year. Or, a customer at a Coast Guard site informs the IM that in the coming quarter that customer will requisition 75 units of an item needed for a special, one-time project. The IM must be able to incorporate such information in the forecast for an item.

The presence of the future quarterly demand buckets allows the IM to do just that. For example, if the IM knows that a planned overhaul in the coming 4th quarter will require 40 units and the forecast does not include that requirement, the IM can add 40 units to the bucket for the 4th quarter.

Multiple-Demand Streams for an Item

The demand for an item may emanate from two or more distinct sources (e.g., shipyard repair and deployed cutters). In such cases, the IM may want to forecast multiple-demand patterns based on the historical demand from each source. (Or, possibly at some future time, future demand for major maintenance activities could be based on parts explosions of end items being overhauled.) If multiple patterns are generated for an item, then its total forecasted demand for a quarter would be the sum of the forecasted demand for that quarter from each pattern.

Procurement Lead-time Forecasting

Procurement lead time affects all items. It is a time period that starts when the need to order stock from a resupplier is established and continues until the stock arrives. Lead-time forecasting is essential to determining when stock should be ordered to avoid a stock shortage. If the forecasted lead time for a procurement is longer than the actual time, stock is being ordered before it is required, more inventory is being held than required, and dollars are being needlessly invested and expended to hold material. If the forecasted lead time is shorter than the actual time, stock may be depleted before a delivery is made, shortages will occur, maintenance will be deferred, and casualties will go uncorrected. A major flaw in many systems is that lead times are typically forecasted

using the most recent occurrence or, at best, the most recent two or three occurrences.

Application

The requirements determination process uses the lead-time forecast for an item to determine its reorder point (i.e., when to order it from a resupplier). Matching that decision with the Ellis and Nathan criteria for selecting a model (see Table 3-2), we have the following:

- *Data* – Few actual observations of lead time for an item are available (unlike demand, which can be based on many observations of an item in a year) and, in the case of that portion of the lead time associated with awarding a contract, differences occur on the basis of the dollar value of the buy.
- *Horizon* – Lead times usually have medium-term time horizons (i.e., 6 months to 2 years) but could have short-term horizons (i.e., less than 6 months).
- *Use* – The reorder point computation is automated or, at least, the initial recommendation to the IM is automated.
- *Cost* – The computation must be made at a low cost without the benefit of the highly specialized expertise of professional forecasters.
- *Single or product line* – It applies to a number of items.
- *Accuracy* – Lead time randomness stems from internal backlog problems and difficulties in awarding orders and from resupplier manufacturing or delivery problems.

Methods, Requirements, and Error Management

Procurement lead time itself is divided into administrative lead time (ALT) – the time from the establishment of the need to the award of an order to a resupplier – and production lead time (PLT) – the time the resupplier takes to fill and deliver the order. We are interested in developing a point estimate for each component so that we can sum them to make our lead-time forecast.[4]

Forecasting Administrative Lead Time

In forecasting ALT, the sparsity of purchase data for an item and the requirement for low-cost forecasts for all items rule out the more sophisticated

[4] Modern inventory management emphasizes the reduction and control of lead times as a means of reducing inventory dollars and being responsive to dramatic changes in demand. We endorse such efforts although they are not the focus of this report.

methods such as decomposition and Box-Jenkins. The sparseness of data also rules out most of the time series models recommended for demand forecasting. A forecast based on the dollar value of the purchase would appear reasonable since the time taken to award a purchase is related to the time taken to prepare a contractual vehicle and method of solicitation. These times depend on the dollar value of the purchase.

To develop an ALT forecast, the Coast Guard could use a fixed industrial engineering standard or a set of standards, but those standards both tend to be too static in reacting to change. We believe that the following model is superior to using one or several standards:

$$Forecasted\,ALT = AVG\,ALT + AVG\,ERR \qquad \text{[Eq. 3-6]}$$

where

$AVG\,ALT$ = the average ALT for all procurements in the last quarter in certain dollar category

$AVG\,ERR$ = the average error for the item between the AVG ALT and the item's actual ALT.

The average error for an item could be computed as follows:

- Every time an item experiences an award, calculate the difference between the actual ALT for that award and the forecasted ALT for dollar value category of the ALT.
- Take the difference and smooth it into a running average using a weight of 0.5 (that weight balances the most recent error and past errors) unless this is the first purchase of the item; if it is the first purchase, use the difference as the average error.

The smoothing computation is given in the following equation:

$$New\,AVG\,ERR = (0.5)(Old\,AVG\,ERR) + (0.5)(Actual\,ALT - AVG\,ALT) \qquad \text{[Eq. 3-7]}$$

This approach for forecasting ALT takes advantage of all recent ALT observations and, at the same time, enables us to consider the continuing differences in ALT that might arise for an item. By including both positive and negative differences in computing the average error term, this approach would yield an error term near zero for an item whose ALTs vary above and below the population average. For an item whose ALTs are consistently above or below the population average, the error term would reflect that difference.

FORECASTING PRODUCTION LEAD TIME

In forecasting PLTs, data are again sparse for individual items. Moreover, PLTs are a function of the vendor and of the contractually required delivery date negotiated with the vendor. Given those circumstances, the last observed PLT for an item or a weighted average PLT for an item with a high weight (say 0.5 or 0.67) would appear appropriate. The latter approach is summarized in the following single exponential smoothing equation:

$$New\,PLT = (1 - Weight)(Old\,PLT) + (Weight)(Last\,Observed\,PLT) \quad \text{[Eq. 3-8]}$$

The problems with this approach are the influence of an outlier (i.e., an abnormal observation) and the failure to consider business cycles. A one-time delivery aberration caused by a strike or a lost shipment could unduly increase the forecasted PLT. To avoid that occurrence, we should use dampening. A reasonable rule might be that any PLT greater than 10 percent or 20 percent above the forecasted PLT would not be automatically smoothed into the forecast unless a review showed that it was not a one-time aberration. (The actual percentage would depend on what each supply center judges as being an acceptable variation.)

The inclusion of business cycles in PLT forecasting is a more difficult problem. Studies show that PLTs are statistically correlated with leading economic indicators. However, no correlation between PLTs and predicted economic indicators has been established. A non-modeling approach entailing market or product analysis could be used to update PLTs. We suggest the Coast Guard use a non-modeling approach for high dollar items and families of technical items that can be identified with a sector of the economy affected by business cycles.

REPAIR LEAD-TIME FORECASTING

For reparable items, repair and not procurement is the principal source of resupply. The Coast Guard relies on several approaches to manage repair. One is to contract for repair and negotiate repair times as part of the award process. Another is to direct reparable units in need of repair (referred to as unserviceables) to a separate maintenance organization that schedules their repair and manages them. Both of these approaches negate the need to forecast repair cycle times (RCTs), but they also introduce the possibility of shortages while unserviceable items are being repaired. A third approach is to include repair times in the requirements process and build or modify inventory levels to consider repair.

Application

To avoid shortages and excesses while unserviceable assets are being repaired, the Coast Guard needs to forecast RCTs accurately prior to actual repair. To do so, the requirements determination process uses an item's RCT

forecast in the same way that it uses the item's procurement lead time to set reorder points. That is, it sets the optimal level of serviceable items below which unserviceable items are sent for repair so that they can be returned in serviceable condition before the on-hand quantity of serviceable items is exhausted.[5]

Repair cycle time comprises several subelements (e.g., induction time, setup time, repair time, inspection time, etc.). Although each of these subelements can contribute to the randomness observed in RCTs, we are only interested in forecasting the total time. Even though RCTs are less than procurement lead times, we would argue that their forecasts are just as critical because reparable items tend to be the most expensive and mission-essential items managed in the inventory. Therefore, if we are overforecasting their RCTs, we are tying up a considerable investment in inventory; and if we are underforecasting their RCTs, we are disrupting the capability of the fleet to perform its missions.

As we discuss in Chapter 5, the requirements determination process uses RCT forecasts to set repair points and safety stocks for reparable items. Matching those applications and RCT traits with the Ellis and Nathan criteria (see Table 3-2) for selecting a forecasting model yields the following:

- *Data* – A degree of randomness would be expected from the variety of potential problems that could be encountered when repairing a unit and the scheduling problems evolving from the need to perform variable workloads with a fixed labor pool.

- *Horizon* – Repair times are less than 1 year and therefore have short- and medium-term forecasting horizons.

- *Use* – The repair point computation is automated or, at least, the initial recommendation to the IM is automated.

- *Cost* – The computation must be made at a low cost without the input of highly specialized expertise of professional forecasters.

- *Single or product line* – It applies to a number of items.

- *Accuracy* – Since forecasted RCTs would go into negotiations with commercial repair sources, as a minimum they need to be realistic.

---

[5]The focus of this section is on forecasting RCTs, that is, on determining the responsiveness of organic and contractor depot maintenance activities to repair orders. Another concern in the area of depot maintenance is the effective use of resources. RCT forecasting does not measure the effectiveness of maintenance resources; it only measures the end product, i.e., how long it takes the assigned resources to get the job done. For a discussion on effective use of maintenance resources, we refer you to the LMI Report AR803R1, *Army Depot Maintenance: More Effective Use of Organic and Contractor Resources*, Kelvin K. Kiebler, Larry S. Klapper, and Donald T. Frank, June 1990.

Methods, Requirements, and Error Management

The generally short-term forecasting horizon for RCTs negates the need for trend, seasonal, and cyclical models and places the focus on levels models or smoothing models. However, given the possibility of a large number of observations, smoothing models would not seem to be a good choice because of the number of repeated computations that would be required. For these reasons, most supply activities rely on either standards or a moving-average model to forecast RCTs.

We believe that the best way to forecast RCTs is to use an 8-quarter, moving-levels model, that is, a levels model that relies on 8 quarters of data on numbers of repair and RCTs to compute an average RCT. Such a model would take advantage of all recent data on RCT while providing for a simple computation.

No modeling of error is recommended. However, as in the case of procurement lead times, we believe that filtering of data is advisable to screen out outliers caused by one-time aberrations. Again, a reasonable rule would be to consider any reported repairs and their associated RCTs that have times greater than some percentage (e.g., 50 percent) of the forecasted RCT for the item before including them in the 8 quarters of historical data. Instead, the IM would review them for inclusion or exclusion. (The actual percentage would depend on what each supply center judges to be an acceptable variation.)

CONCLUSIONS AND *RECOMMENDATIONS*

Demand Forecasting

Demand forecasting is at the center of the requirements determination process, affecting the computation of all stockage rules. The randomness of demand causes any forecast to be wrong. The objective is to reduce the forecast error by including in the forecast as much information as possible on the historical demand patterns, future program data for items related to an equipment system, and any specific information an IM might receive on future demand.

We recommend that the Coast Guard

- *Collect historical demand for every item in two sets of 8 quarterly buckets – one for recurring demand and one for nonrecurring demand – and that forecasting be done on a quarterly basis.*

- *For items having demand in 2 or more of the past 8 quarters, compute a demand forecast as follows:*

  - Starting with the two sets of historical demand, develop (1) for low-priced items, a historical demand pattern summing recurring and non-recurring demand and then dampening quarterly values to three standard deviations or (2) for high-priced items, two historical demand patterns, one only recurring demand and the other total demand, and then dampen to 2 standard deviations.

  - Use the demand pattern(s) to compute and store forecasts for the 11 models in Table 3-5 not associated with program data.

  - For items associated with a program, use the actual 8 quarters of demand and program data to compute a usage rate and compute and store forecasts for the three models in Table 3-5 associated with program data.

  - For the 11 or 14 models that apply, use stored demand and forecast data to compute MSEs.

  - Use the MSEs to select the model with the lowest error.

  - Use that model to construct a future demand pattern for a minimum of 8 quarters.

Procurement Lead-Time Forecasting

Lead-time forecasting is essential to the proper timing of orders to replenish stock. Procurement lead time is composed of the time from the establishment of the need to the award of an order to a resupplier, that is, the ALT, and the time the resupplier takes to fill and deliver the order, that is, the PLT.

We recommend that the Coast Guard

- *Forecast the ALT for an item as the sum of the following:*

  - The average ALT for all procurements in the last quarter in the dollar value category into which the item's order quantity falls

- The smoothed error between the population ALT and any actual ALTs for the item.

- *Forecast the PLT for an item using an exponential smoothing model with a high weighting constant[6] and filtered lead-time data (i.e., only observed times that are representative and not one-time aberrations).*

Repair Lead Times

To avoid shortages or excesses while unserviceable assets are being repaired, the requirements determination process should forecast RCTs prior to actual repair.

We recommend that the Coast Guard

- *Maintain for each reparable item historical data on RCTs by collecting the past 8 quarters of repairs and associated times that are filtered (i.e., only observed times that are representative and not one-time aberrations)*

- *Forecast the RCT for a reparable item using an 8-quarter moving average model.*

[6] A weighting constant of 1 would be equivalent to using the last observed PLT.

CHAPTER 4

Item Management

INTRODUCTION

Several alternatives exist for providing supply support for the population of items managed by the Coast Guard. Some items warrant supply-center-level stockage and others do not because their low demand and their low essentiality relative to the missions of the Coast Guard. Once customer demand for an item is predicted, the first inventory management decision that must be made is how best to provide supply support. Not only must IMs chose between stocking or not stocking an item, they must also choose between several possible alternatives.

Alternatives

Within the Federal regulations governing acquisition and management of materiel, the following alternatives are generally available for managing an item:

- Local purchase
- Centrally managed, nonstocked, procure to fill demand
- Centrally managed, nonstocked, order from contract to fill demand
- Centrally managed, stocked with traditional demand-supported levels
- Centrally managed, stocked with insurance or low-demand levels.

While each alternative has a cost and a degree of responsiveness associated with it, an item's characteristics may rule out consideration of a particular alternative. Therefore, to be consistent with our overall requirements determination objective of providing a high level of customer support at the lowest cost, the selection of an item management alternative should be based on costs, customer support requirements, and the type of item (see Figure 4-1).

LOCAL PURCHASE

Under the local purchase alternative, the customer does not normally requisition an item from the supply center but instead fills a demand by buying the item locally. The availability of the commercial product locally is a prerequisite

for this alternative. Items with special specifications or items that are sole source purchases are not good candidates for local purchase.

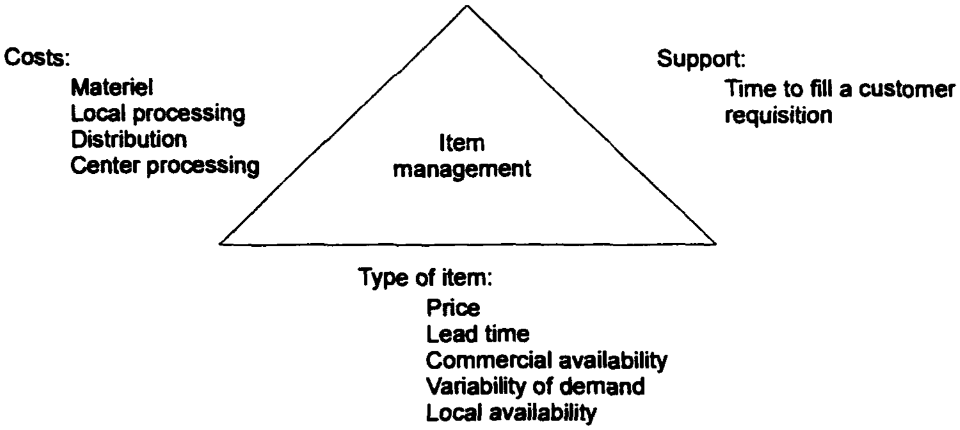

Figure 4-1.
General Factors Involved in Item Management

Centrally Managed, Nonstocked, Procure to Fill Demand

Under this alternative, each customer requisition is filled by a procurement action. Since procurement actions have a cost associated with them, this alternative is normally reserved for items with little or no predicted demand. Items with long lead times that are normally needed by the customer in a short time are not good candidates for this alternative.

Centrally Managed, Nonstocked, Order from Contract to Fill Demand

Under this alternative, a call-type contract is established for an item or group of items and each customer requisition is filled by an order placed on that contract. This alternative can take on many forms. One form is a centrally negotiated contract but decentralized ordering because the contract involves a supply schedule or catalog that is given to customers for direct ordering.

Another form made possible by today's technologies is a centrally negotiated contract with centralized ordering; under that arrangement, the supply center electronically transmits customer requisitions to the vendor for direct delivery. This alternative entails the Coast Guard entering into a call-type contract that establishes an agreement between the vendor and Coast Guard on pricing and responsiveness. The vendor and not the Coast Guard maintains stock. Government agencies using this alternative have encountered instances in which a vendor may have a minimum order quantity for shipping materiel. To overcome that restriction, they have either contracted with a second vendor to handle

requisitions for small quantities or elected to stock some amount of the item to fill requisitions with quantities below the vendor's minimum.

This alternative requires that the item have enough demand (sales) to warrant the cost of establishing a call-type contract and to ensure that commercial vendors are willing to enter into this type of arrangement. Commercial items are good candidates for this alternative as are items with limited shelf lives or hazardous items that require special storage since stockage has additional costs and risks for these items. The grouping of similar candidate items for a larger sales volume on one contract increases the potential willingness of vendors to participate while spreading the cost of the contract across many items.

Centrally Managed, Stocked with Traditional Demand-Supported Levels

Under this alternative, the supply center maintains stock to fill customer requisitions. The level of stock is a function of the item's demand and price and the customer responsiveness goal. This alternative is applicable to any item that has sufficient demand to economically support stockage and that cannot be more economically supported through a call-type contract.

Centrally Managed, Stocked with Insurance or Low-Demand Levels

Under this alternative, a supply center maintains stock to fill customer requisitions. Although the objective is to maintain a minimal level of stock, the actual quantity depends on the price of the item and the customer responsiveness goal. This alternative is applicable to any item with a low demand but a high responsiveness requirement. It would also be applicable to an item with no predicted demand but one for which any backordered demand arising from catastrophic failure (e.g., fire) would prevent the unit from performing its mission.

General Approach

Historically, to determine whether an item should be stocked, supply activities have relied on manual analysis of individual items, range rules, or depth rules (zero depth means no stockage). Manual analysis is the most subjective approach and also the most costly since it is labor-intensive and time-consuming. It is the least desirable approach.

For most items, range rules are the preferred method for determining how to manage an item. In theory, those rules can include economic and non-economic considerations and can address all potential alternatives for managing an item. In practice, they are traditionally formulated as either demand frequency rules or an economic trade-off model that only compares the estimated cost of stocking an item against the estimated cost of not stocking it. In either case, the model is

often referred to as a range model since it determines the range of items that are to be stocked.

A growing trend, particularly with the advent of readiness-based sparing models in the provisioning of new end items or major systems, is to use a depth model (i.e., a model to calculate how much to stock) to determine whether to stock an item. If the depth model computes a zero inventory level for an item, it is not stocked. Using depth rules to determine stockage is appealing, particularly when dealing with items that directly affect the operation of end items. However, for other items, depth rules are less appealing because they normally do not treat all costs included in item management. Rather, they focus only on the costs of the material.

Limitations on Migration Between Alternatives

The following conditions limit the movement of an item from one management alternative to another:

- Once an item is designated for local purchase, the supply center loses visibility of demands. Therefore, it is difficult, if not impossible, to change its status. An exception might occur if the supply center learns that a product is not available at the local level either from customers, from a large number of local demands referred to the supply center for acquisition, or from a reporting process that accesses local purchase files.

- If demand for a mission stopper (an item whose failure precludes mission performance) were to decline to zero, the item would still be stocked as an insurance item.

- If abundant assets are on hand for an item, it cannot be a candidate for commercial support until those assets are exhausted or reduced to a level that can economically be carried during the length of the contracted commercial support.

- If an item is under an established call-type contract, such as a requirements contract, it cannot be switched to a stocked item until the contract expires or is terminated. (A contract could be terminated for reasons such as noncompliance or quality problems. These problems are outside the scope of selection modeling.)

Selection Model

As previously noted, the selection of an management alternative for an item depends on cost, support, and item factors. A selection model logically combines those factors into a decision algorithm, or in this case, a decision tree. To define

the decision tree that is appropriate for the Coast Guard, we need to consider the three categories of factors.

Cost Factors

The costs involved in the stockage decision span most of the costs involved in inventory management. They include the administrative costs in managing and procuring an item, the distribution costs in warehousing and transporting materiel, and the cost of the materiel itself.

Table 4-1 lists the variable cost factors associated with each item management alternative. To include them in the selection model, we need to place them in a total variable-cost equation. The equation serves as a mathematical laboratory that will give the cost of managing an item under an alternative for a given period of time (a year is a generally acceptable period). We can then rank alternatives according to their costs. For a given item, the least-cost alternative is the likely candidate for selection within the boundaries set by the support and item factors for that item. Although we listed cost factors for local purchase, we contend that the selection of this alternative is totally dependent on local product availability as defined under item factors.

When developing values for these costs, we are only concerned with the variable or differential portion of each cost. For example, even though the cost of the supply center commander is part of the cost of centrally managing items, that cost should not be included in the administrative cost of centrally managing an item because it will not change no matter which alternative is selected for an item.

Definition of Costs

In this subsection, we define the variable cost factors listed in Table 4-1. We also indicate whether the cost is treated as a one-time or recurring cost in its associated total cost equation(s).

The *administrative cost of introducing an item* into central management is the sum of the costs of establishing technical and supply data records for a new item and of assigning a management team (i.e., an IM and others as required). Since this one-time cost is common to all alternatives except local purchase, it is excluded from their total-cost equations.

The *administrative cost of centrally managing an item* is the sum of the costs of maintaining item technical and supply data records and maintaining a management team. Since this annual recurring cost is common to all alternatives except for local purchase, it is also excluded from their total-cost equations.

When a customer demands a nonstocked item from a supply center, the supply center must initiate a procurement to fill that demand. The *cost of making a*

Table 4-1.
Item Management Cost Factors

Item management alternative	Variable cost factors
Local purchase	1. Cost of making a local purchase 2. Local purchase price of an item 3. Time-weighted cost of a backorder
Nonstocked, procure to fill demand	1. Administrative cost of introducing an item *(new item)* 2. Administrative cost of centrally managing an item 3. Cost of making a procurement 4. Central procurement price of item with direct delivery to the customer 5. Backorder cost of a lead-time delay
Nonstocked, order from contract to fill demand	1. Administrative cost of introducing an item *(new item)* 2. Administrative cost of centrally managing an item 3. Cost of establishing call-type contract 4. Cost of placing an order against the contract 5. Contract price of an item with direct delivery to customer 6. Time-weighted cost of a backorder
Stocked, demand-supported levels	1. Administrative cost of introducing an item *(new item)* 2. Administrative cost of centrally managing an item 3. Cost of initial stock *(new item)* 4. Cost of carrying stock 5. Cost of ordering stock 6. Procurement price of the item 7. Cost of filling a requisition/demand 8. Time-weighted cost of a backorder
Stocked, insurance or low-demand levels	1. Administrative cost of introducing an item *(new item)* 2. Administrative cost of centrally managing an item 3. Cost of initial stock *(new item)* 4. Cost of carrying stock 5. Cost of ordering stock 6. Procurement price of item 7. Cost of filling a requisition/demand 8. Time-weighted cost of a backorder

procurement is the sum of the administrative costs involved in awarding a contract and in providing any postaward contract administration. The *central procurement price of item with direct delivery to the customer* is the price that the supply center pays for the material plus any premium that might be incurred because of

a low volume, special delivery, or a restrictive customer-required delivery date. (For purposes of cost comparisons between alternatives, we are only interested in the premium, i.e., any price difference from the normal procurement price for the item if such a difference exists.) Both of these costs are recurring ones depending on the number of demands and their quantities.

The *cost of establishing call-type contract* is the sum of the administrative costs involved in awarding an indefinite-type contract and in providing any post-award contract administration. This is a one-time cost, which, if the contract covers a group of items, is reduced by the number of items in the group. The *cost of placing an order against the contract* is the sum of the administrative costs involved in ordering from an indefinite-type contract. This is a recurring costs that depends on the number of demands. The *contract price of an item with direct delivery to customer* is the negotiated price for material ordered under the contract. (Again, for purposes of cost comparisons between alternatives, we are only interested in the premium, i.e., any price difference from the normal procurement price for the item if such a difference exists.) This is a recurring cost that depends on the demand quantities filled through the contract.

For a stocked item, inventory can be divided between operating stock (i.e., stock that is procured, issued, and then procured again) and safety stock (i.e., stock that is procured but not normally issued except to avoid backorders). When the decision is made to stock a new item or an item not previously stocked, the *cost of initial stock* is the one-time material cost of safety stock.

The *cost of carrying stock* is the sum of the cost of capital (the cost of money – normally set at 10 percent in the Government), the cost of storage, the cost of obsolescence, and the costs of pilferage, spoilage, and damage. This is an annually recurring cost normally expressed as a percentage of the dollar value of the average inventory on hand.

The *cost of ordering stock* is the sum of the administrative costs of procurement, the warehousing costs of receiving and stowing material, and if not included in the price of the material, the transportation cost of moving the material from the vendor to a storage location. This is a recurring cost dependent on the number of procurements placed in a year. The *procurement price of the item* is the price paid for the item and is a recurring cost dependent on the amount of material ordered in a year.

The *cost of filling a requisition/demand* is the sum of the administrative costs of processing a requisition at a supply center, the warehousing costs of issuing the item, and the transportation costs of distributing the material from a storage location to the requesting unit.

The *time-weighted cost of a backorder* is a unit-of-time cost for a backorder. It includes record keeping and emergency procedure costs as well as any operational costs that are incurred awaiting the item. The latter is particularly difficult to establish, especially when the backorder hinders mission performance. (In such cases, the cost of the next higher assembly may represent the cost of the

backorder.) Time weighting places emphasis on the responsiveness of an alternative to a delayed issue. This is an annual recurring cost dependent on the expected number of backorders on hand at any point in time for an alternative.

Cost analysis cannot handle every reason for stocking or not stocking an item. For example, quantifying the degree to which an item is readily available at the local level is difficult if not impossible. Hence, guidance on stockage policies relies on human judgment and often uses terms such as "must consider," "must balance," and "should include" when discussing noneconomic criteria. However, we believe the use of noneconomic criteria must be kept to a minimum and strictly disciplined. To do otherwise would allow subjectivity to take over and defeat the goal of high support at the lowest cost, which is inherent in the analytical approach to selecting an item management alternative.

Total Variable Cost Equations

For the alternatives other than local purchase, we offer the following total variable cost (TVC) equations (each represents the expected annual costs of the alternative):

- Nonstocked, procure to fill demand

TVC = total costs of procurement + premium cost of material + backorder costs

or mathematically,

$$TVC = A_1 r + d(P_1 c) + b_c(rl) \qquad \text{[Eq. 4-1]}$$

where

r = annual number of requisitions for item

A_1 = cost of making a procurement

d = annual demand for item in units

P_1 = any premium paid for material (10 percent = 0.1)

c = procurement price of item

l = lead time for item less average time if item stocked

b_c = time-weighted backorder cost for item.

The total cost of procurement is the cost of a procurement multiplied by the number of procurements which is, in this case, the number of requisitions. The premium cost of material is the number of units ordered (the forecasted demand

for the item) multiplied by the premium, which is given as a percent of the price of the item. The backorder cost is the unit-of-time cost for a backorder multiplied by the number of backorder units of time, which are the products of the number of requisitions and the added lead time for not stocking the item. The total variable cost is the sum of these costs.

- Nonstocked, order from contract to fill demand

TVC = cost of contract + total cost of orders + total cost of material + backorder costs

or mathematically,

$$TVC = A_2 + Or + d(P_2c) + b_c(l_cr) \qquad \text{[Eq. 4-2]}$$

where

r = annual number of requisitions for item

A_2 = cost of establishing the contract

O = cost of placing an order

d = annual demand for item in units

P_2 = any premium paid for material

c = procurement price of item

l_c = lead time to order and deliver from contract minus time for item stocked

b_c = time-weighted backorder cost for item.

The cost of contract is the cost of establishing the contract against which orders are placed. The total cost of orders is the cost of an order multiplied by the number of orders, which is the number of requisitions. The premium cost of material is the number of units ordered (the forecasted demand for the item) multiplied by the premium, which is given as a percentage of the price of the item. The backorder cost is the unit-of-time cost for a backorder multiplied by the number of backorder units of time, which is the product of the number of requisitions and the added lead time for not stocking the item. The total variable cost is the sum of these costs.

- Stocked, either stockage alternative.

TVC = new item costs + total costs of ordering + total costs of carrying + backorder costs

or mathematically,

$$TVC = \min[0, c(s-a)] + A_3\left(\frac{d}{q}\right) + (H \bullet c)\left(\frac{q}{2} + s\right) + b_c b_o \qquad \text{[Eq. 4-3]}$$

where

min [x,y]	=	minimum of x or y
c	=	procurement price of item
s	=	safety stock level under this alternative
a	=	current assets for item (0 for new item and items not previously stocked)
A_3	=	cost to order
d	=	annual demand for item in units
q	=	order quantity for item under alternative
H	=	holding cost rate
b_o	=	on-hand backorders times 365 days in year
b_c	=	backorder cost for item.

The new item cost is the cost of bringing the minimum level of stock on hand to the level of safety stocks for the item. The total cost of ordering is the cost of a procurement multiplied by the number of procurements, which is the annual demand quantity for the item divided by quantity per procurement. The total cost of carrying inventory is the cost of carrying a unit, which is the product of the holding rate and the item's unit price, multiplied by the average quantity on hand, which is one-half of the order quantity plus the level of safety stocks. The backorder cost is the unit-of-time cost for a backorder multiplied by the number of backorder units of time. The total variable cost is the sum of these costs.

Although backorder costs are formulated for all alternatives as the backorder-unit-of-time cost multiplied by the number of backordered units, the number of backordered units is different for each management alternative. For the procure-to-fill-demand alternative, each requisition is "backordered" and the time to fill is the differential lead time (i.e., the normal procurement lead time less delivery and backorder time if the item were stocked). For the order-from-

contract alternative, each requisition is again "backordered" and the time to fill is the differential order time (i.e., the normal time to order and deliver material from the contract less delivery and backorder time if the item were stocked). The order lead time would be subject to negotiation with the vendor. For purposes of alternative selection, the differential order time could be assumed to be equal to zero or to the time for the stocked, demand-supported alternative (i.e., the contractor experiences the same backorders as organic stockage).

For the stocked, demand-supported alternative, the number of backordered units depends on the item's safety stocks. For purposes of alternative selection, the system's responsiveness goal and average time on backorder could be used to compute that number as follows:

$$b_o = d(b_t - r) \qquad \text{[Eq. 4-4]}$$

where

b_o = on-hand backorders

d = demand for item

b_t = system average time on backorder

r = system response time.

The same number could be used for the stocked, insurance/low-demand alternative.

Once system-level costs [shown in capital letters in Equations 4-1, 4-2, and 4-3 (e.g., A_3 = cost to order)] are developed for the alternatives, specific item parameters (shown in small letters) can be fed into the equations and the low-cost alternative selected. Alternatively, the equations could be used to develop tables that would reflect the low-cost alternative for different combinations and ranges of item parameters. However, with this approach the IM cannot see the differences between alternatives.

Support Factors

Responsiveness and essentiality are two critical, noneconomic support factors that must be considered when selecting an item management alternative. If a unit cannot wait an extended period of time for a demanded item because it would impede the performance of the unit's mission, certain alternatives may need to be ruled out.

Since local availability is a prerequisite for local purchase, support is not an issue for this alternative. However, if a unit that has any indication that a

problem might exist with local availability, it should reconsider the local purchase selection since that alternative would not be responsive.

Responsiveness translates to a delay equal to one lead time under the nonstocked, procure-to-fill-demand alternative. If that delay is unacceptable because of the essentiality of an item to Coast Guard units, this alternative must be ruled out. If a vendor cannot guarantee an acceptable response time under the nonstocked, order-from-contract alternative and once again the item is essential to Coast Guard units, that alternative must be ruled out.

Responsiveness is normally not a concern for the stocked alternatives except in setting safety stocks. It is a principal concern in decisions related to the stockage alternative such as the decision to position inventory forward and disperse it rather than hold it in a central location.

Item Factors

Other noneconomic factors that affect item management selection are the item's characteristics. For instance, local availability of an item is a prerequisite to local purchase. (We distinguish between being able to buy locally and being able to buy from across the country and around the world through a local buyer. Local purchase means the former not the latter.) Items that are not commercial products, items that are only available sole source, or items whose unit cost exceeds the purchase ceiling are generally not locally available. Also, items with long lead times (i.e., greater than a year) might not be available locally.

Long lead times (greater than 1 year) should also rule out items from the non-stocked, procure-to-fill-demand alternative since the response time would be unacceptable to any unit.

Items identified as commercial products, not having unique specifications, would be candidates for the nonstocked, order-from-contract alternative. Other items would need to be reviewed by buyers to determine if this alternative could meet their requirements, that is, to determine if a contract could be awarded for these items. To avoid the extra costs associated with losses from expired usage dates or special handling requirements, the Coast Guard may wish to emphasize contractibility and seek out vendors for items with limited shelf life or hazardous storage requirements.

Whether or not items either are locally available or a contract can be awarded for them entails some subjectivity. To help decide on either approach, an expert system could be used to improve the quality of human judgement and to guide less experienced personnel in making decisions. Through interviews with subject matter experts, the key item characteristics could be identified and weighted. With the elements and their weights, an expert system could be developed that would guide IMs through a complete item characteristics analysis.

CONCLUSIONS AND *RECOMMENDATIONS*

The five general alternatives for managing an item are as follows:

- Local purchase
- Centrally managed, nonstocked, procure to fill demand
- Centrally managed, nonstocked, order from contract to fill demand
- Centrally managed, stocked with traditional demand-supported levels
- Centrally managed, stocked with insurance or low-demand levels.

The selection of an item management alternative should be based on costs, responsiveness or support to the customer, and the type of item.

We recommend that the Cost Guard adopt the model shown in Figure 4-2 for selecting an item management alternative.

The following apply to the certain branching questions in the decision tree:

- As previously noted, the determination of local availability requires an assessment of several item characteristics.
- If an item is on an established contract but that contract will expire within a lead time, it should not be considered "under contract" in the decision tree (see Figure 4-2).
- The determination of whether an item can be placed under a call-type contract requires procurement and source of supply expertise.
- When assets are "abundant" depends on the demand rate for the item and the inventory carrying costs associated with the assets.

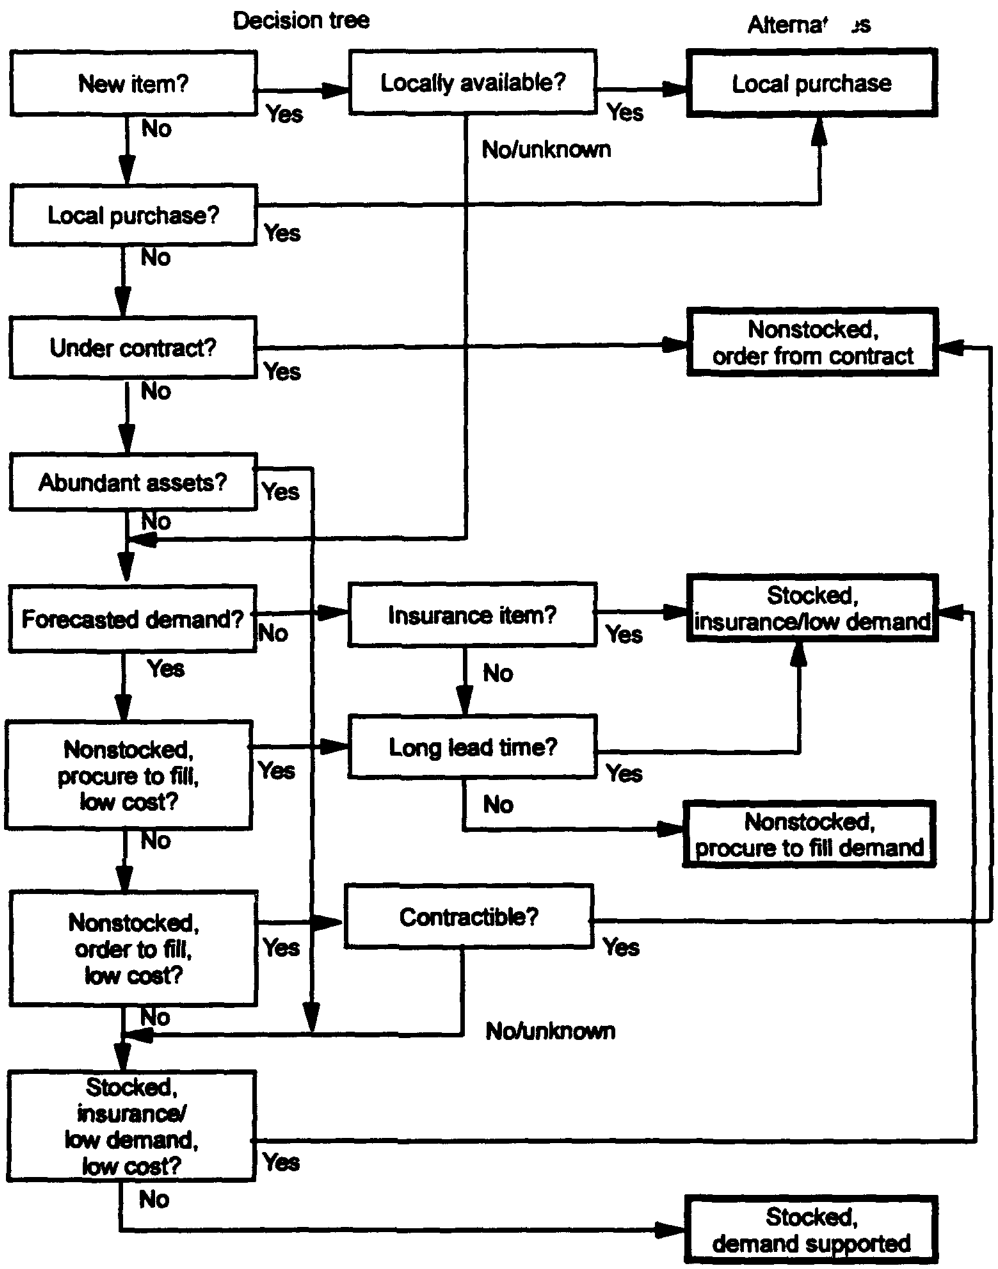

Figure 4-2.
Proposed Model for Item Management Selection

CHAPTER 5

Demand-Supported Stockage

INTRODUCTION

If a supply center decides to stock an item, the next decision is how much to stock. The answer depends on whether the item is stocked with demand-supported levels or with insurance or low-demand levels. We discuss demand-supported stockage in this chapter and insurance and low-demand stockage in Chapter 6.

As we previously noted, the Coast Guard's supply center inventory system does not revolve around scheduled or dependent demand. Rather, it maintains inventory to fill random or independent demand from a variety of Coast Guard units. As demands for an individual item draw down its inventory level, the system orders stock from a source of supply (either a distributor or manufacturer) to replenish that level. If the demand for an item is stable (a level demand pattern), this process of repeating stock depletion and replenishment generates the classical inventory saw-tooth curve shown in Figure 5-1.

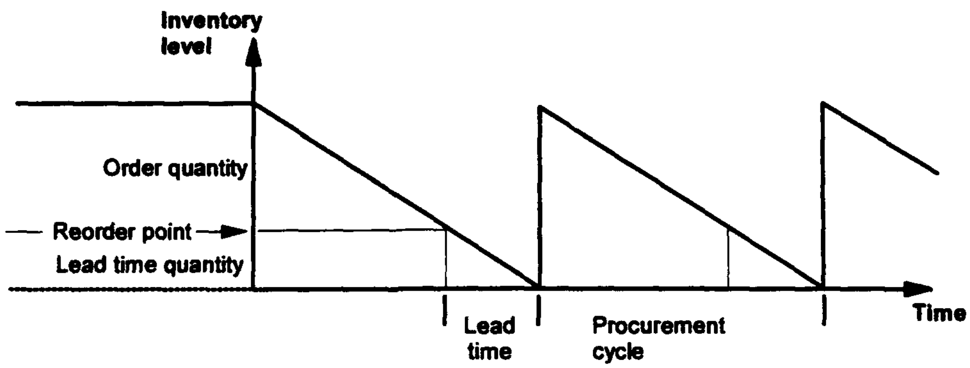

Figure 5-1.
Classical Saw-tooth Curve for On-Hand Inventory

As illustrated, this process requires that the system know when to order (i.e., have a reorder point) and how much to order (i.e., have an order quantity). If an item has demand-supported stockage, the answers to these two questions are built on the forecasted demand for the item.

Demand-Supported Levels

To establish the quantities of operating and safety stocks that it has on hand and on order, an independent demand inventory system has traditionally adopted the approach of constructing the series of levels as shown in Figure 5-2.

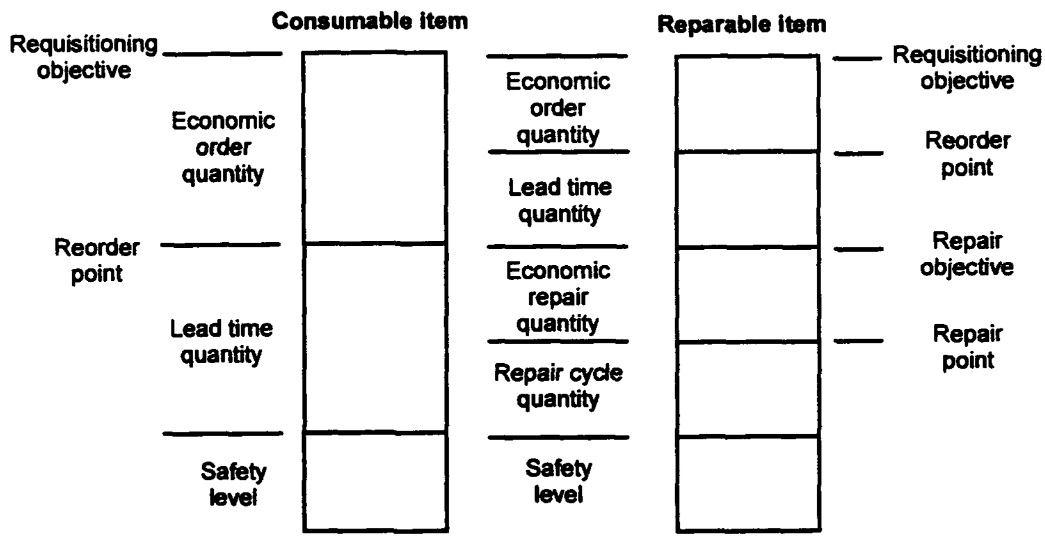

Figure 5-2.
Traditional Levels of Inventory

Definitions

The following definitions apply to the levels:

Requisitioning objective – When demands draw down an item's inventory and an order is placed to replenish that inventory, the system orders up to the item's requisitioning objective (or ordering objective).

Economic order quantity – The economic order quantity (EOQ) is the amount of stock in a replenishment order. (In this term, "economic" means the quantity represents the low-cost solution to the problem of how much to buy.)

Reorder point – The reorder point (or order point) is the level of inventory at which a replenishment order is initiated.

Lead Time Quantity – The lead-time quantity is the amount of stock required to meet demand during the lead time (i.e., ALT and PLT) of a replenishment order.

Safety Level – The safety level is a quantity of stock maintained to protect against backorders during the period of a replenishment order (and during the period of a repair order for reparables).

Repair Objective – When demands draw down an item's serviceable inventory and a repair order is placed to maintenance to replenish that inventory, the system directs repair of assets up to the item's repair objective.

Economic Repair Quantity – The economic repair quantity is the amount of stock in a repair order. (In this term, "economic" means the quantity represents the low-cost solution to the problem of how much to repair.)

Repair Point – The repair point is the level at which a repair order is initiated.

Repair Cycle Quantity – The repair cycle quantity is the amount of stock that is required to meet demand during the repair cycle of a repair order.

INVENTORY

Levels Interaction

For a consumable item, the system uses the inventory levels as follows:

- As shown in Figure 5-2, an item's reorder point is the sum of its lead-time quantity and safety level. When the item's inventory position (i.e., the sum of its on-hand and on-order assets minus any quantity on backorder) is at or below the reorder point, the system places an order.
- The size of that order is the difference between the requisitioning objective and inventory position. As shown, if the inventory position is at the reorder point, that difference would equal the economic order quantity. If demand varies and the inventory position actually drops below the reorder point, the size of the order is the sum of the economic order quantity and the reorder point deficit (i.e., the difference between the reorder point and the inventory position). However, when the order is awarded to a vendor, the actual order quantity may be adjusted to meet the vendor's minimum-buy quantity or to take advantage of a quantity discount.
- Since demand and lead times will vary from the time a replenishment order is initiated until it arrives, a safety level quantity is needed to protect against stockouts or backorders.

For reparable items, the situation is complicated by the fact that they have two sources of supply – procurement and maintenance, that reparable assets can be either serviceable (ready for issue) or unserviceable (not ready for issue, in need of repair), and that a demand for a serviceable unit is not always accompanied by the return of an unserviceable unit. (An unserviceable unit may not

always be returned because the demand may be for an initial issue, a shipment may be lost, or a customer may have failed to ship the unserviceable unit.).

- A reparable item's reorder point is the sum of its PLT quantity and its repair objective; and when the item's inventory position is at or below the reorder point, a procurement order is placed.

- The procurement-lead-time quantity is based on that portion of demand not filled by repairing unserviceables.

- The size of the procurement order is the difference between the requisitioning objective and inventory position and is equal to the economic order quantity. However, the order quantity only replenishes that portion of demand not satisfied by a repaired unserviceable item (i.e., if an unserviceable item is condemned or if no unserviceable item is turned in with the demand for a repaired item).

- The item's repair point is the sum of its repair lead-time quantity (normally referred to as its repair cycle quantity) and its safety level; when the level of serviceable items on hand and in repair is at or below the repair point, unserviceable items are sent to maintenance for repair.

- The size of the repair order is the difference between the repair objective and the level of serviceable items on hand or in repair and is equal to the economic repair quantity. However, the repair quantity only replenishes that portion of demand that is reparable.

- Since demand and procurement and repair lead times vary, the safety level protects against backorders until either serviceable items arrive from maintenance or a replenishment order arrives.

In summary, the computation of demand-supported levels for consumable items involves order quantities; reorder points; lead-time quantities; safety levels; and, for reparable items, it additionally involves repair quantities, repair time quantities, and repair points.

Total Variable Cost Approach – Theory and Practice

In theory, the problem of determining when to order and how much to order is approached either as a "Q,r problem," i.e., a problem whose solution is the order quantity, Q, and the reorder point quantity, r, or an "S,s problem," i.e., a problem whose solution is the point when you order, s, and the point you order to S. (For reparable items, an S,s problem is often approached as an S,S-1 problem because the order quantity is equal to one.) Although the form of the problem is different between Q,r and S,s problems, they are really equivalent since s equals r and S equals the sum of Q plus r. We focus on the Q,r form of the problem.

The theoretical approach to solving the Q,r problem is to develop a total variable cost equation and solve for the Q and r that minimize that equation. The most basic TVC equation is:

TVC = total cost of ordering + total holding cost

or mathematically,

$$TVC = AD/Q + (HC)(Q/2 + R - L) \qquad \text{[Eq. 5-1]}$$

where

A = cost to order

D = annual demand for item (in units)

Q = order quantity

H = holding cost rate

C = procurement price of item

R = reorder point

L = lead-time demand quantity.

This equation is similar to Equation 4-3, the TVC equation for item management stockage, except that it excludes the new item costs and backorder costs. The total cost of ordering is the cost of a procurement multiplied by the number of procurements, which is the annual demand quantity for the item divided by quantity per procurement. The total cost of carrying inventory is the cost of carrying a unit, which is the product of the holding rate and the item's unit price, multiplied by the average quantity on-hand, which is one-half of the order quantity plus the reorder point minus the lead-time demand quantity (the difference is the level of safety stocks). (The new item costs are excluded since they are not a consideration in determining how much to stock but rather only when to stock. To avoid duplicate terms, backorder costs are excluded at this point since some form of backorder constraint is normally placed on the equation and is mathematically treated as an additional backorder cost expression to the equation.)

If we were to solve Equation 5-1, we would find that Q is equal to the classical economic order quantity and r is the lead-time demand quantity. However, the problem is normally not as simple as shown in Equation 5-1 because performance constraints are added to the equation. Those constraints direct the solution toward achieving a performance goal and are mathematically expressed as inequalities (e.g., greater than or equal to, less than or equal to).

Some examples of constraints would be as follows:

- Supply availability performance goal – This constraint places a lower limit on supply availability or the percentage of demand that is filled immediately (in the inequality, immediate issues are expressed as the total demands minus the number of backorders established).

$$\frac{Total\,demand - Backorders\,established}{Total\,demand} \geq Goal \qquad \text{[Eq. 5-2]}$$

Since total demand is a given constant in the inequality, by varying the supply availability goal, we constrain the maximum number of backorders that are established.

- Responsiveness performance goal – This constraint places an upper limit on the mean system response time given in Equation 2-2. In Equation 2-2, if the time to make immediate issue is zero, the mean system response time can be shown to be mathematically equivalent to the backorders on-hand divided by the total demand.

$$\frac{Backorders\,on\,hand}{Daily\,total\,demand\,rate} \leq Goal \qquad \text{[Eq. 5-3]}$$

Since the number of backorders on-hand is the product of the number of backorders established and the time on backorder and the daily total demand rate is again a given constant, we constrain the combination of the number and duration of backorders by varying the response time goal.

In theory, Q and r are found by solving the constrained equation for Q and r that yields the minimum TVC. However, practitioners have found it simpler to find Q with a standard economic quantity formula, put that solution into the constrained equation, and solve for r. Since that approach reduces the complexity of the problem and provides a good approximation to the optimal solution, we believe it is a suitable procedure for the Coast Guard.

Periodic Versus Continuous Review and Levels Computation

An important timing issue in applying demand-based levels is how often assets are checked against the reorder point and, also, how often reparable items are checked against the repair point. Under continuous review, the inventory system performs the check whenever assets or levels change. Under periodic review, the inventory system performs the check between intervals of time (e.g., each month). Daily processing is considered continuous review and not periodic review.

Periodic review has the advantage of less checking but has the disadvantage of requiring additional stock to account for any possible delay in placing an order while waiting to perform a check. Today's computer power diminishes the

need for periodic review or for large numbers of items, reduces the periodicity to days, which is small enough to be considered continuous review. *We recommend that the Coast Guard perform continuous review.*

Another important timing issue is the frequency of computing levels. Under continuous computation, the system recomputes levels whenever a change in assets or in a factor that goes into the computations (e.g., lead time changes or price changes) occurs. Under periodic computation, the system recomputes levels at regularly scheduled points in time (e.g., monthly, quarterly).

In Chapter 2, we recommend that the Coast Guard forecast demand on a quarterly basis to avoid interjecting too much randomness in its forecasts. Given that demand and forecast error are the major factors in computing levels, we believe that the computation of levels should track with the forecasting of demand, i.e., the Coast Guard should also recompute levels on a quarterly basis.

Single-Versus-Multiple-Item Modeling

When used to refer to the computation of an inventory level, the term "single item" means that the level for an item is computed independently of levels for other items. For example, the classical economic order quantity model is a single-item model in that it computes the order quantity for an item independent of other items.

In contrast, the terms "multiple item," "multi-item," or "item group" refer to an item-level computation that is not independent of the other items in the group. For example, the safety level that a system-safety-level model computes for items is aimed at meeting a system-wide supply performance goal. (As used in this report, item group does not refer to a family of items that are interchangeable or that can be substituted one for another. We use the term item to refer to a bachelor item, i.e., one that cannot be interchanged or substituted, or the head of a family.)

Multi-item management is not new to the Coast Guard because supply centers establish blanket purchase agreements (BPAs) and basic ordering agreements (BOAs) for groups of items. Although this acquisition practice improves materiel management by reducing lead times and the need for procurement resources and by ensuring higher product quality and more responsive vendor support, it is not related to multi-item inventory models. An item under either a BPA or BOA can have its demand-supported levels computed using either a single-item model or a multi-item model.

Multi-item inventory models are usually identified with greater efficiency because they tend to make tradeoffs across items and reap the benefits of economies of scale. However, multi-item models do not exist for all requirements determination decisions. Where they do exist, they may not be applicable because the lack of special group data or because no logical grouping exists.

We address both single- and multi-item models later in this chapter.

Single-Echelon Versus Multiple-Echelon Modeling

Within the computation of levels, the term "single echelon" refers to an approach that considers the supply activity as a single, autonomous level of supply. As illustrated in Figure 2-2, customers place demands on the activity, the activity fills as many of the demands as possible with issues from its stock and establishes backorders for the rest to be issued when stocks are replenished. The activity replenishes its stocks with orders placed on its sources of supply.

In contrast, the terms "multiple echelon" and "multi-echelon" refer to an approach in which a chain of two or more supply activities is considered. Figure 5-3 illustrates a two-echelon system. Customers place demands on the first activity in the chain and that activity fills as many of the demands as possible and either backorders or refers the rest to the next activity. The first activity replenishes its stocks with requisitions placed on the next activity in the chain, and that activity, unless it is the last activity in the chain, fills as many of the requisitions and referrals as possible and either backorders or refers the rest to the next activity. The last activity in the chain fills both requisitions and referrals from lower echelons and replenishes stocks from its sources of supply. An activity in the chain may choose not to stock an item and refer all requisitions for that item to the next higher echelon or in the case of the last activity in the chain, the source of supply.

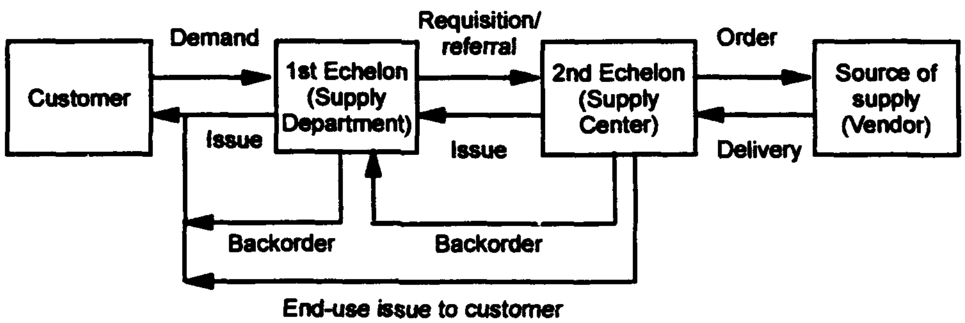

Figure 5-3.
Two-Echelon Supply Support

In principle, when more than one echelon of supply supports a customer, a multi-echelon model should compute levels that are as effective but cost less than if each echelon has its levels computed with single-echelon models. However, the multi-echelon model requires the centralized computation of levels (normally at the highest echelon) and knowledge of demand at the lowest echelon by the centralized computational activity. Most important, the validity of the

model's computations is directly proportional to the accuracy of the data and forecasts that go into it.

Our analysis of its supply operations shows that the Coast Guard must take the following actions before it can apply a multi-echelon inventory model with centralized computation of levels:

- The supply centers must change business practices and compute both their levels and the levels for supply activities that directly support Coast Guard operational units.
- The supply centers must know the demands that units place on all levels of supply activities. If they are to have that knowledge, data must be collected at a variety of sites and transmitted to the supply centers.
- The Coast Guard must perform a statistical analysis of the data that would go into a multi-echelon computation (e.g., distributions for site demand and for order-and-shipping times between sites and the supply center). Without such an analysis, the multi-echelon model has to be built on standard assumptions, which may not be accurate. In the end, without the analysis the multi-echelon computation could produce supply levels that cost more and provide less effective support.
- The application of a model that simultaneously computes levels for multiple echelons must currently be limited, for the most part, to high-cost reparable items that operate in an S,S-1 environment (see page 5-5). Items operating in a Q,r environment must rely on single-echelon models since a viable multi-echelon Q,r model still needs to be developed.[1]

In summary, our discussion focuses on single-echelon models, not because we believe that multi-echelon models offer no benefit to the Coast Guard but because we believe that a number of obstacles exist that hinder their successful application to the Coast Guard requirements determination process.

Economic Order Quantity

The computation of the economic order quantity (EOQ), or lot-sizing as it referred to in the literature, answers the question of how much to order. In general, Government supply activities use the classical Wilson EOQ model, or some variation of it, to determine how much to order. (The exceptions occur mostly with reparable items and insurance and low-demand items.)

The classical EOQ model trades off the cost of ordering stock against the cost of holding stock to determine the best order quantity. It assumes level (constant)

[1] Multi-echelon models currently in use are (S,S-1) or (S,s) models, which are not concerned with economic order quantities. Some heuristic methods have been suggested for handling low-cost items for which the batching of replenishment orders is a key economic concern; however, none of those methods are currently being used.

demand, a single cost to hold that is proportional to the value of the inventory, and a single cost to order. Table 5-1 shows the relationship between the size of the EOQ and an item's demand and unit price. Activities using the model generally place maximum and minimum constraints on the quantity generated by the model. (For example, the Department of Defense restricts wholesale order quantities to a maximum of 24 months of demand and a minimum of the lesser of 6 months or the administrative lead time.)

Table 5-1.
Demand, Unit Price, and EOQ Size

Item attribute	Relationship to EOQ		
	Conversational	Logical	Mathematical
Forecasted demand (D)	Directly proportional	If D $\uparrow$ then EOQ $\uparrow$ and if D $\downarrow$ then EOQ $\downarrow$	Percent change in EOQ = Square root of percent change in demand
Unit price (C)	Inversely proportional	If C $\uparrow$ then EOQ $\downarrow$ and if C $\downarrow$ then EOQ $\uparrow$	Percent change in EOQ = One over square root of percent change in price

In the private sector, *dependent demand* found in the manufacturing process has sparked the development of several lot-size models that have fixed planning horizons and nonlevel-demand patterns. These models use the same costs as the classical EOQ models with their independent demand, but they rely on a heuristic approach rather than an equation to determine the optimal quantity. Except for the mathematical relationships, Table 5-1 also applies to nonlevel EOQs.

Our recommendation to forecast a future demand pattern versus a single-point estimate opens the possibilities of both level and nonlevel demand. Therefore, both categories of EOQ models have potential application in the Coast Guard.

The Simplest Order Quantity Models

Before we discuss models for level demand and models for nonlevel demand, we need to say a few words about the simplest models for computing order quantities.

Lot-for-Lot

The lot-for-lot (LFL or L4L) model is the most elementary and straightforward lot-sizing technique. The order quantity is simply set equal to the

forecasted demand for the next period of time (e.g., the forecasted demand for the next quarter). Since the order quantity is consumed in the period in which it arrives, minimal holding costs are incurred. However, the annual ordering costs are equal to the number of periods (orders) in a year times the cost to order. This technique is best used when demand is known and when the ordering cost is low (e.g., when ordering on of an established contract) and carrying costs are high (e.g., a high-priced item with relatively small and irregular demand).

Although the theory of order quantities has long passed the LFL model, the model is experiencing new success in the field of manufacturing. Here, lead times are short (e.g., 2 weeks or less), demand is known with some certainty, and the cost of ordering is low because the manufacturer can transmit needs to an established supplier by electronic data interchange. Under these circumstances, the LFL model's simplicity and ease of use make it an attractive alternative to the more complex models. However, since these circumstances do not exist at a Coast Guard supply center, we believe this model is not appropriate for Coast Guard use since it does not trade off costs to arrive at the least cost solution.

Fixed-Period Model

Simplicity and ease of use make the fixed-period, or days-of-supply, model (e.g., 6-month order quantity for all items) an appealing option. However, like the LFL model (which is a special case of this model), the fixed EOQ model is not really an EOQ model since it does not trade off costs in arriving at the order quantity. We do not believe it is appropriate for the Coast Guard.

Level-Demand Models

The following models are EOQ models that assume that future demand is steady. Each one is tailored to a specific situation but offers potential benefits for the Coast Guard.

Classical or Wilson Economic Order Quantity

The EOQ is the best-known and most frequently applied lot-sizing technique. It seeks to minimize the total costs of holding and ordering inventory. Variables in the EOQ model are the annual unit demand for an item, D; the cost-to-order, A; the holding cost rate, H; the unit price for the item, C; and, of course, the order quantity, Q. By using those variables, the model can formulate a TVC equation and solve it for Q. In formulating the TVC equation, the model is deterministic, that is, it assumes that demand and costs are known and relatively stable.

The TVC equation and EOQ formula are as follows:

$$TVC = A\frac{D}{Q} + HC\left(\frac{Q}{2}\right) \qquad \text{[Eq. 5-4]}$$

$$Q = \sqrt{\frac{2AD}{HC}} \qquad \text{[Eq. 5-5]}$$

The derivation of the EOQ formula from the TVC equation is taken from the theory of differential calculus and can be found in an inventory textbook such as Hadley and Whitin, *Analysis of Inventory Systems.*[2]

Quantity Discount Model — Total Units Variation

The whole, or all units, discount EOQ model is a variation of the classical EOQ. It handles discounting where Unit Price 1 (C_1) applies for order quantities in the interval from 1 to Breakpoint 1 (P_1); Unit Price 2 (C_2) applies for order quantities in the interval from 1 to breakpoint 2 (P_2); etc. For example, if the unit price for order quantities up to 100 is $10 per unit and for order quantities over 100 is $8 per unit, then 90 units would cost $900, while 110 units would cost $880.

For a price interval, the TVC equation and optimal Q formula are as follows:

$$TVC_j = DC_j + A\left(\frac{D}{Q_j}\right) + HC_j\frac{Q_j}{2} \qquad \text{[Eq. 5-6]}$$

$$Q_j = \sqrt{\frac{2AD}{HC_j}} \qquad \text{[Eq. 5-7]}$$

The model's optimization procedure is to compute the optimal Q for the first interval. If that Q is not within the interval endpoints (and therefore not possible), the optimal Q for the interval is the breakpoint that yields the lowest cost in the interval total cost equation. That process is repeated for all price break intervals. The overall EOQ is the interval Q among all the interval Q's that yields the lowest total cost (TVC_j).

Quantity Discount Model — Incremental Discount Variation

The incremental, or partial, discount EOQ model handles discounting where Unit Price 1 (C_1) applies for units in the interval from 1 to Breakpoint 1 (P_1); Unit Price 2 (C_2) applies to units in the interval from Breakpoint 1 (P_1) plus 1 to Breakpoint 2 (P_2); etc. The total cost of *n* units where *n* is in the interval *j* defined by Breakpoints $P_j - 1$ and P_j is the cost of the units up to $P_j - 1$ (R_j) plus the price in the interval times the number of units in the interval (i.e., $n - P_j - 1$). For example, if the unit price is $10 per unit for up to 100 units and $8 per unit for units

---

[2] Hadley, G., and T.M. Whitin, *Analysis of Inventory Systems*, Prentice-Hall, Inc., Edgewood Cliffs, N.J., 1963.

starting at 100 and above, then 90 units would cost $900 while 110 units would cost $1080.

For a price interval, the TVC equation and EOQ formula are as follows:

$$TVC_j = DC_j + \frac{D}{Q_j}(A + R_j - C_jP_{j-1}) + \frac{HR_j}{2} + \frac{HC_j}{2}(Q_j - P_{j-1}) \qquad \text{[Eq. 5-8]}$$

$$Q_j = \sqrt{\frac{2D(A + R_j - C_jP_{j-1})}{HC_j}} \qquad \text{[Eq. 5-9]}$$

The model's optimization procedure is the same as the all-units discount procedure. That is, it computes the optimal Q for the first interval. If that Q is not within the interval endpoints (and therefore not possible), the optimal Q for the interval is the breakpoint that yields the lowest cost in the interval total-cost equation. This procedure is repeated for all price break intervals. The overall EOQ is the interval Q minus all the interval Qs that produces the lowest total cost.

Group Economic Order Quantity

The group EOQ model addresses the case in which all items in a group are procured at one time instead of one item at a time. The rationale behind this approach is that it will reduce the overall cost to order since the cost of procuring a number of similar items at one time is not n times the cost to order (A) but is A plus n – 1 times the cost of an additional item on an order (a). The formula for the group EOQ in dollars ($Q) is:

$$\$Q = \sqrt{\frac{2(A + (n-1)a)\sum_{i=1}^{n} D_iC_i}{H}} \qquad \text{[Eq. 5-10]}$$

Once the group dollar EOQ is found, individual item order quantities are computed on the assumption that each item's order quantity is a multiple of that item's annual demand (D) and that the multiple is the same for all items in the group. The equations for solving for the multiple m and for each item's order quantity, Q_i are:

$$m = \frac{\$Q}{\sum_{i=1}^{n} D_iC_i} \qquad \text{[Eq. 5-11]}$$

$$Q_i = mD_i \qquad \text{[Eq. 5-12]}$$

Multiple Cost-to-Order Model

The multiple cost-to-order model addresses the cases in which an item could experience different ordering costs depending on the method of procurement (e.g., small-purchase versus large-purchase procedures). For each method applicable to the item, the model uses that method's ordering cost in the classical EOQ formula and TVC equation to generate an order quantity and cost for that method. If the optimal order quantity is outside the range for that method of procurement (e.g., if the optimal Q for a small-purchase ordering cost has a dollar value greater than the small-purchase limit), then the quantity is revised to the closest size to the optimal quantity within the range and the TVC equation is used to develop a revised cost for the method. The costs for applicable methods are compared and the order quantity is set to the quantity for the method with the lowest cost.

Economic Order Quantity Model with Space-Based Storage Cost

The space-based storage cost model treats bulky, low-cost items whose storage costs are not proportional to their prices. The model separates the holding rate in the classical EOQ model into all storage cost for an item and a nonstorage cost holding rate, H_{ns}. The storage cost is formulated as the expected space occupied (i.e., volume, v) multiplied by the cost of space, c_s. The EOQ formula is:

$$Q = \sqrt{\frac{2AD}{2vc_s + H_{ns}C}} \qquad \text{[Eq. 5-13]}$$

Nonlevel-Demand Models

We use the term "nonlevel-demand models" to describe the dependent-demand[3], lot-size algorithms we believe can be used to handle nonlevel-demand patterns in the Coast Guard. Recently, with the emphasis on material requirements planning in manufacturing, a number of these types of models have arisen. In what follows, we list the most recognized of these algorithms.[4] Since the mechanics of these techniques cannot be summarized in an equation, Appendix C gives a detailed explanation and an example of each algorithm.

- *Least total cost* (LTC) — Using the standard cost to order and cost to hold, the LTC algorithm computes the total costs for all possible options for buying demand over a fixed period of time and then selects the option with the

---

[3] Dependent demand is associated with a known demand process such as the manufacturing of an end item. For example, if an automobile manufacturing plant was going to build 500 cars in the coming month and each car had four door handles, then the plant has a dependent demand forecast of 2,000 handles for the coming month.

[4] Because of the great interest in identifying the best-ordering practices for manufacturing, the literature contains a number of papers on proposed new lot-size techniques for independent demand. We limited our list to those that are commonly found in material requirements planning software.

lowest total cost. Alternatively, it computes the carrying cost for each buy and compares it to the cost to order and, like the classical EOQ, selects the buy in which they are closest to being equal. This approach could become increasingly time-consuming as the number of periods being considered is extended.

- *Part-period balancing* (PPB) – "Part-periods" refer to the number of periods that parts are held in the inventory. Developed by IBM, the PPB technique starts by computing a critical value equal to the cost to order divided by the cost to hold one period. Then, for buy quantities beginning with the demand for the first period and adding demand period by period, it computes the part-periods for each buy as long as the cumulative number of part-periods is less than the critical value. When the part-periods exceed the critical value, PPB chooses the buy quantity whose part-periods are closest to the critical value.

- *PPB with look-ahead/look-back tests* – Like the other algorithms in this category except for LTC, PPB is a technique that is designed to quickly search out the best buy among all possible buys. However, like the others, it may sometimes select a good buy but not the best. To improve the effectiveness of the technique, IBM added two tests. The look-ahead test checks to see whether the number of part-periods is lower when the next period is added to the current buy or is used as the first period in the next buy. The look-back test checks to see whether the number of part-periods is lower when the last period in the current buy is used as the first period in the next buy.

- *Group PPB* – The group PPB combines the PPB technique with the savings derived from a group purchase. It uses an ordering cost equal to the standard ordering cost plus the incremental cost increases for additional items on the contract. It then rolls the demand for all items in the group into a group demand pattern. The normal PPB technique relies on the quotient of the ordering and carrying costs as a critical value for selecting the best buy quantity. Since the carrying costs will vary by period when dealing with a group of items with different costs and demand, the group PPB uses the ordering cost as the critical value and compares it to the part-period carrying costs of buys for increasing periods of demand.

- *Least unit cost* (LUC) – The LUC algorithm uses the standard cost to order and cost to hold to calculate the lot size as the number of periods of demand that have the smallest per unit cost.[5] Starting with the first period, it totals the costs for buying one period of demand and divides by the number of units in the buy to arrive at unit cost for a one-period buy. It repeats the computation for a two-period buy, a three-period buy, etc. It stops at the buy when the unit cost stops decreasing and when the next buy with an additional period has a larger unit cost. This algorithm has been found to be erratic in identifying low-cost solutions.

---

[5] Unit cost in this case does *not* refer to the price of the item. It is the total system costs of a purchase spread over the units in that purchase.

- *Wagner-Whitin (W-W) heuristic approach* – The W-W heuristic approach is a dynamic programming technique for determining the low-cost lot size. It starts by comparing the cost of one order for the first two periods against the cost of two orders, one for each period and selecting the low-cost solution. Then, it moves to three periods, deciding between one order for all three periods or the low-cost solution for two periods plus one order for the additional period. The calculation continues in this fashion until N periods are completed, where N periods is the preset time horizon. Because of its complexity, practitioners have refrained from using the heuristic.

- *Silver-Meal (S-M) heuristic approach* – The S-M technique is another somewhat complex heuristic approach for determining the low-cost lot size. First, we compute the total cost per unit of a unit buy; that is, a buy that covers the first period of demand. Then, compute the total cost per unit of a two-time unit buy; that is, a buy that covers the first two periods, and continue until the cost per unit stops declining. The optimal lot size is equal to the demand over the number of periods for which the lowest cost per unit of time occurred. Like the W-W heuristic approach, this approach has not seen wide-spread application.

Periodicity of the Nonlevel-Demand Model

The nonlevel-demand models were developed for use in a manufacturing environment in which monthly or weekly demand varies but is known. We believe such models have the potential to outperform EOQ models (which deal best with steady demand patterns) for items with quarterly future demand patterns that vary up and down. Although using these models for future demand patterns that vary seems reasonable, it does introduce the additional cost of buying in terms of units in a period instead of individual units. (For example, if the forecasted demand for two periods is 10 units per period, the choice between buying 10 or 20 units may not produce the least-total-cost solution when it is 15 units.)

Arguably, the error introduced in estimating the costs that go into a model could be larger than the error from period buying. However, since we recommend quarterly forecasting and not weekly or monthly forecasting, we feel the additional cost from the quarterly buy is too large to ignore. To reduce that cost, we believe that, for a nonlevel-demand model, demand forecasting should still be done quarterly, but the forecasted quarter should be broken down into months.

If the forecasted quarter is based on a seasonal or usage forecasting model, divide the forecast by 3 to arrive at a monthly forecast. If, example, the forecast is 20, assign 7 to months one and two and 6 to month three – assign remainders to early months to hedge against shortages. If the forecasted quarter is based on a trend model, develop the three monthly forecasts using the trend model.

The current dependent-demand models do not directly address quantity discounts or the possibility of multiple costs to order. However, their inclusion in the different nonlevel-demand models is not as difficult as in the case of the level-demand models where an iterative approach is required. In that case, when the quantity increases to a discount or order-cost breakpoint, the next lot size under consideration would use the new unit price or order cost, respectively.

Proposed Model

We believe the decision tree shown in Figure 5-4 is suitable for selecting a model to compute an order quantity for an item.

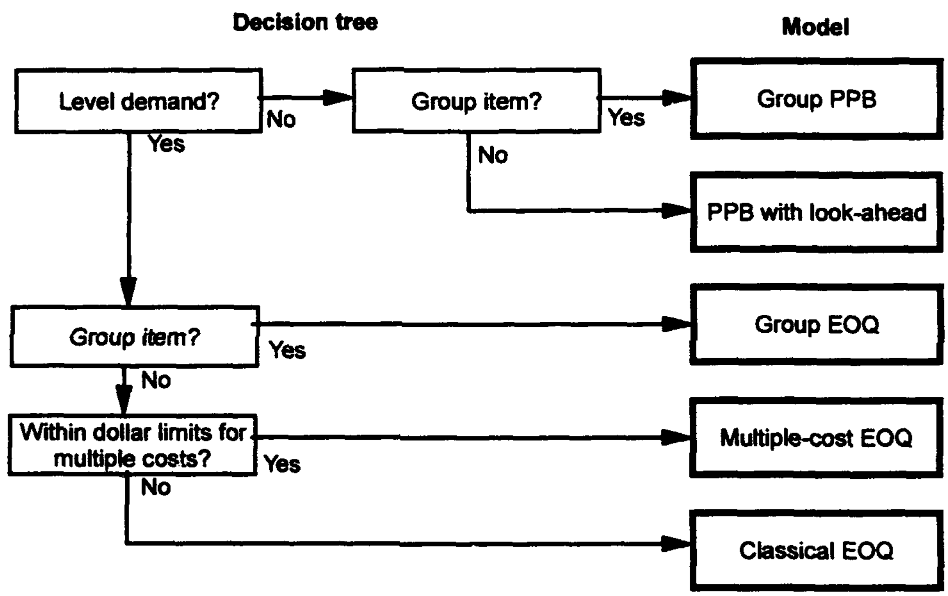

Figure 5-4.
Proposed Methods for Setting Order Quantities

The following conditions apply to the certain branching questions in the decision tree:

- The assignment of an item to a group is independent of the order quantity computation.
- The quantity discount models are not part of the illustrated proposal since they come into play during negotiations at the time of award. If the vendor

offers a discount during these negotiations, the models could be used to determine whether a quantity adjustment is justified.

The computed order quantity derived from this decision tree goes into the setting of the requisitioning objective and, as is subsequently discussed, the setting of the safety-level quantity. The actual recommended buy quantity when the reorder point is breached is the difference between the requisitioning objective and the inventory position at that time (which should be the same quantity as the computed order quantity if the inventory position equals the reorder point). The recommended buy quantity could be modified to account for quantity discounts or minimum vendor buy quantities.

Reorder Point

When the inventory position for an item (i.e., the level of stock on hand and on order) drops below its reorder point, it is a signal for an order to be placed. The reorder point for a consumable item is normally computed as the sum of an item's lead-time quantity, its safety level, and any special levels. That is,

Reorder point = lead time quantity + safety level + special levels

or mathematically,

$$r = (L_q)(D_q) + s + l_s \qquad \text{[Eq. 5-14]}$$

where

r = reorder point

L_q = lead time (in quarters)

D_q = demand (forecasted quarterly demand)

s = safety level

l_s = special level.

We do not discuss special level except to say that stock can be reserved for special programs. For example, assets might be retained for predetermined specific conditions or contingencies such as to support an emergency relief or rescue mission.[6] In theory, the stock in special levels should be protected by operating and safety stocks until needed.

[6] Special levels should not be established for the expected demand for a future known or scheduled maintenance plan because such demand should be more correctly included in the forecasted demand pattern for the item.

For a reparable item, the reorder point is the sum of the repair cycle quantity, the lead-time quantity, the safety level, and any special requirements for special projects.

Reorder point = lead time quantity + repair cycle quantity + safety level + special levels

or mathematically,

$$r = (L_q)(D_q - G_q) + (R_q)(G_q) + s + l_s \qquad \text{[Eq. 5-15]}$$

where

r = reorder point

L_q = lead time (in quarters)

D_q = demand (forecasted quarterly demand)

G_q = regenerations (forecasted quarterly repairs)

R_q = repair cycle time (average time to repair in quarters).

In Equation 5-15, we introduced the variable "regenerations." It pertains to those unserviceable items that are turned in for a serviceable item and are subsequently repaired. The differences between the regenerations and demand are first, unserviceable items that are turned in for a serviceable item but are found to be beyond repair and are condemned and second, demands for which no unserviceable item is turned-in (e.g., a new demand or an unserviceable item being lost in shipment).

The number of regenerations can be derived either by forecasting them in a manner similar to forecasting demand or by tracking the average portion of demands that is regenerations (i.e., demands that are accompanied by returned unserviceable items that can be repaired rather than condemned). In practice, although the former sounds more tedious, it may be simpler because it does not require the matching of returning unserviceable items with demand for serviceable items. It does require the tabulation of the ratio of returned unserviceable items to demands and the percentage of unserviceable items sent to maintenance that are repaired.

Procurement Lead-Time Quantity

As discussed in Chapter 3, procurement, or acquisition,[7] lead time is the time between initiation of an order and reception of the ordered material ready for issue. As such, it includes the administrative lead time needed to award an

[7] The popularity of the term "acquisition lead time" rather than "procurement lead time" is increasing, particularly with DoD and other Government suppliers.

order, the production (i.e., manufacture or stock pick, and pack) lead time and the transportation lead time needed to deliver the order. To preclude any stock outage, an order must be placed one procurement lead time before the material is actually needed. The procurement lead-time quantity (PLTQ) is the quantity of an item that will be requested during the period equal to the lead time. As shown in the sawtooth curve in Figure 5-1, the reorder point is not a level of stock used but rather a breakpoint used to determine when an order needs to be placed for the item.

As shown in the reorder point equations, an item's PLTQ is computed by multiplying its demand and its procurement lead time. For a reparable item, the demand used in the computation is the demand not satisfied through regeneration. This computation is correct when the forecasted demand pattern is level.

However, the very nature of reparables often causes nonlevel-demand patterns that we must accommodate in the PLTQ computation. The result is that the PLTQ is not a constant at any point in time as it is with level-demand procedure; rather, it will vary depending on where we are in the varying demand pattern. Consequently, we start our quarterly PLTQ computation by determining where we are on the pattern and when we expect to place an order. Then, we continue forward in the pattern to determine the demand that we expect over the item's lead time.

We illustrate this nonlevel-demand procedure with an item that has the following forecasted demand pattern:

Quarter	1	2	3	4	5	6	7	8
Forecasted	10	2	5	10	10	2	5	10

If the item has 5 units on hand and zero on order, we anticipate that we have enough units to cover the first 1/2 quarter of demand. If the item has a lead time of 6 months, we convert the 6 months into 2 quarters and divide that 2 over the quarters of forecasted demand as follows:

2 quarters = 1/2 of 1st quarter + 2nd quarter + 1/2 of 3rd quarter.

PLTQ = 5 units + 2 units + 3 units = 10 units.

If the same item has 15 units on hand and zero on order, then we anticipate covering the first 2 3/5 quarters of demand. We then divide our 2 quarters as follows:

2 quarters = 2/5 of 3rd quarter + 4th quarter + 3/5 of 5th quarter.

PLTQ = 2 units + 10 units + 6 units = 18 units.

If the item has 5 units on hand and 10 units on order but due in at the end of the first quarter, we still compute the PLTQ as above. However, we would anticipate

that we would have backorders between the halfway point in our first quarter (when on-hand stock is exhausted) until the on-order stock arrives. Releasing another procurement would not help this situation because that order would take 2 quarters to arrive. Expediting the current outstanding order is the only alternative for possibly avoiding a backorder situation.

If the item's inventory position is such that the next expected order is past the 8 future quarters of forecasted demand, we compute the item's PLTQ using the level-demand procedure (i.e., multiply the item's demand in quarters by its lead time in quarters to arrive at the PLTQ). If our demand is a level 10 units per quarter, applying the nonlevel-demand procedure yields the same answer as the level-demand procedure, namely, 20 units.

Safety-Level Quantity

In computing order quantities, lead-time quantities, and reparable repair cycle quantities, demand and times are assumed constant, or in statistical terms, "deterministic." In reality, they are not constant; they vary, or in statistical terms, they are stochastic. To guard against shortages due to varying demand and lead times, inventory systems have safety-level quantities. As illustrated in Figure 5-5, safety level stocks fill increases in demand or lead times that would otherwise deplete stock.

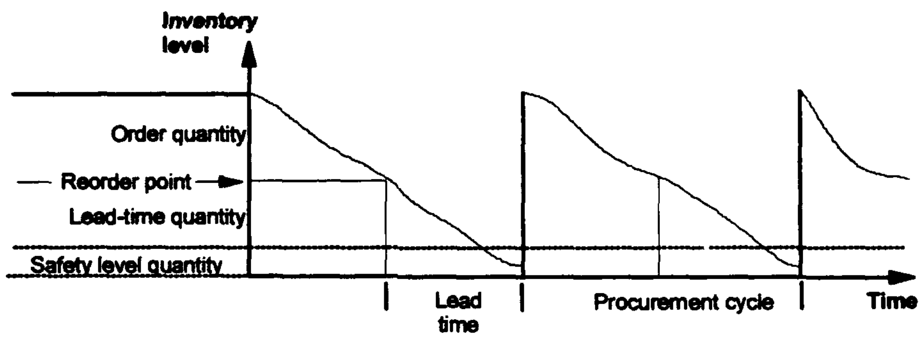

Figure 5-5.
On-Hand Inventory Curve for Stochastic Demand

Basic Theory

We have referenced statistical terminology in defining the safety level. That referencing is appropriate since the safety-level formulation is well grounded in statistics. The theory is that demand is a variable that can be described with a statistical distribution, such as a normal, or bell shaped, distribution or a Poisson distribution for items with low demand. The distribution gives the probabilities

for different levels of demand occurring and is referred to as a probability density function.

Two statistics of importance to any statistical distribution are its mean and its standard deviation. The mean of the lead-time demand distribution is the expected lead-time quantity. That is, when we stock the lead-time demand quantity, we have enough stock to cover all probabilities of demand during a lead time up to the mean. For example, if the lead-time demand were 10 units, we could cover demand of zero through 10; however, if the actual lead-time demand exceeded 10, we would have backorders. To avoid those backorders, we need a safety-level quantity to cover the spread of demand greater than the mean.

Standard deviation (=) is a measure of the spread of a distribution. If a safety level is to provide a given level of protection against backorders, that is, to cover a given percentage of possible demand probabilities, it must cover a certain amount of the spread of the probability function. Consequently, an item's safety level is mathematically defined as a multiple of the standard deviation of the probability density function describing the distribution of units demanded during the lead time:

$$s = k\sigma \qquad \text{[Eq. 5-16]}$$

where

s = safety level

k = multiple

σ = standard deviation of probability density function.

Computing the Standard Deviation

An exact calculation of the standard deviation is not possible since we do not know the exact probability density function for an item; rather, we rely on a function we believe is close. A precise calculation can be made by using the observed demand history as a sample and then computing the sample standard deviation as follows:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n} (d_i - d_{av})^2}{n-1}} \qquad \text{[Eq. 5-17]}$$

where

d_i = ith observed demand for a lead time

d_{av} = average of all lead-time demand observations

n = number of observations.

In practice, the mean absolute deviation over a lead time (MADLT) is used in place of this computation because the relationship between the MADLT and σ is approximately

$$\sigma = 1.25\,MADLT \qquad \text{[Eq. 5-18]}$$

The following points govern the use and computation of the mean absolute deviation (MAD):

- When setting safety levels, we are not really interested in how demand varies; rather, we are interested in how well we are doing in forecasting demand. If we could reduce our forecast error to zero, that is, accurately predict demand no matter how it varies, we would not need safety-level quantities to protect against demand variance. Indeed, studies have shown that safety levels built on forecast error outperform safety levels built on demand variance. That finding makes sense since we are using our forecasts to determine what stock we need. The level of error in our forecasts is what we want to guard against.

- For computing the MAD, either the straight average or the smoothed average is acceptable:

$$MAD = \frac{\sum_{i=1}^{n} |D_i - F_i|}{n} \qquad \text{[Eq. 5-19]}$$

$$New\,MAD = (Old\,MAD)(1 - weight) + (Forecast - actual)(Weight) \qquad \text{[Eq. 5-20]}$$

- We are not interested in the variance of demand over a quarter (our forecast interval); we are interested in the variance of demand over a lead time. Therefore, we must adjust the MAD to reflect MADLT. In practice, a power rule[8] such as the one that follows with a β factor of 0.7 gives reasonably good results:

$$MADLT = MAD\left(\frac{L}{F}\right)^{\beta} \qquad \text{[Eq. 5-21]}$$

[8] Within DoD, a number of mathematical expressions are applied by the different Services and the Defense Logistics Agency (DLA). Most of the expressions may be found in, *Stockage Policy Analysis*, Office of the Assistant Secretary of Defense (Manpower. Reserve Affairs, and Logistics), Final Report, Annex A, Part 3, 31 August 1980, pp. 1 - 12 to 1 - 14.

where

L = lead time in months

F = forecast interval (i.e., 3).

- This MADLT formulation only considers demand variability and not lead-time variability. (This is somewhat standard since in most cases we do not have enough lead-time observations to develop an accurate standard deviation.)

Early Modeling Approaches

When safety levels were first developed, one approach to setting them was to rely on the engineer who designed the item. That engineer knew the item's failure rate and could use probabilities to estimate what was required to prevent any shortage. The problem with this approach was that the actual environment that generated the failures and the subsequent demand was different or changed from the environment on which the engineer based the estimate. Furthermore, the engineer was concerned with the variance of the item's demand and not with the error in our forecasting demand.

Another early modeling approach was the fixed safety level; that is, safety level quantities set equal to a certain number of days of supply. The number of days was usually based on some estimate of the average variance in demand for the population of items. The problems with this approach were that it did not consider the differences in demand and lead-time variances between items, it tended to be static (i.e., once the number was set it was rarely changed), and it did not consider the costs in determining the size of safety levels.

Fill-Rate or Supply-Availability Models

The fill-rate models evolved from the days-of-supply models. They were developed to satisfy a demand fill rate or supply availability for an item. (Sometimes, they are referred to as requisitions, short, or confidence-level models.)

The first versions of these models were single-item models; that is, they attempted to give the same level of protection to every item. Although they overcame the problems of differences in demand variances between items and static computations, they still had the fixed-safety-level problem of not considering costs.

The next versions of these models were multi-item models; that is, they attempted to provide a level of protection across all items. They traded off the cost of preventing a backorder for a high-cost item against the cost of preventing several backorders for low-cost items. Presutti and Trepp present two versions of

a multi-item, supply availability safety level model.[9] They used the Method of Lagrange Multipliers to solve for k_i, the safety-level factor for item i.

The Presutti and Trepp models address units backordered. To convert their unit models to requisition models, that is, to protect against frequency or number of backorders instead of quantity backordered, one divides the demand in the models by the item's average requisition size.

Response-Time Modeling

In their paper, Presutti and Trepp also gave two versions of a multi-item, response-time, safety-level model.[10] They refer to such models as time-weighted models. These safety models differ from their supply-availability counterparts in the way they consider backorders. The supply-availability models focus on the expected number of backorders established (EBEs), while response-time models focus on the expected number of backorders on hand at any random point in time (EBOs). The relationship between the two is as follows:

$$EBO = (EBE)(B_t) \qquad \text{[Eq. 5-22]}$$

where

B_t = average time on backorder.

That is, the backorders on hand equal the number established per day multiplied by the days on backorder. The difference between the two types of models is that the response-time model also considers duration of a backorder and duration is important to the customer whose requisition is placed on backorder.

Safety-level Constraints

Most inventory systems place minimum and/or maximum constraints on safety levels. A typical minimum might be zero and a typical maximum might be a lead time's demand or three standard deviations of lead-time demand. The former constraint falls out of the fact that managers have difficulty with the concept of negative safety levels or a negative level of any kind for that matter. The maximum constraints are set to restrict investment in a single item and to limit reliance on the number coming out of the tail of the probability density function (a distribution is valid only up to a point).

---

[9] Presutti, Victor J., Jr., and Richard C. Trepp, "More Ado About Economic Order Quantities (EOQ)," *Naval Research Logistics Quarterly*, Vol. 17, No. 2, June 1970, pp. 243-251.

[10] Ibid.

Proposed Model

We believe the multi-item, response-time, safety-level model is most suitable because

- It considers the differences in forecast error between items.
- It is not a static computation.
- It trades off costs among items.
- It considers both the number and duration of backorders.

The mathematical form of the model is from Presutti and Trepp's Model IV[11] with the substitution of 1.25 MADLT for σ. The precise form is as follows:

$$s_i = 1.25\, k_i\, MADLT_i \qquad \text{[Eq. 5-23]}$$

where

s_i = safety level for item i

$MADLT_i$ = MADLT for item i

and

$$k_i = 0.707 \ln \left[\frac{MADLT_i SYSSUM(1 - \exp(-1.13 Q_i / MADLT_i))}{2.56 Z_i Q_i C_i \beta} \right] \qquad \text{[Eq. 5-24]}$$

where

k_i = safety-level factor for item i

Z_i = average requisition size for item i

Q_i = order quantity for item i

c_i = price for item i

β = system backorder target

$SYSSUM$ = $o\ [c_i\, MADLT_i]$ for all items

ln = natural logarithm function

exp = exponential function.

---

[11] Ibid.

Like any safety-level model, this model has some shortcomings. It does not consider the following real-world conditions that would affect performance:

- Management actions such as backorder reconciliation and cancellations, direct deliveries, expediting, and lateral resupply from other Coast Guard units could reduce backorder time.
- A portion of the item population will have assets in excess of their requirements.
- Demand may not be normally distributed as assumed in this model.
- Our proposed computation of *MADLT* is centered on demand variance and does not directly consider lead-time variance. (Failure to consider that variance does not mean the mathematics is too difficult; rather, the sparsity of lead-time data for an individual item precludes any accurate measure of lead-time variance.)[12]
- Maximum and minimum constraints could be placed on safety levels.

With regard to this last point, we believe that the Coast Guard should adopt a maximum safety level of three standard deviations of lead-time demand and a minimum of one negative standard deviation. The idea of a negative level does not usually make sense, but in this case it has merit. The procurement-lead-time quantity is based on the expected demand for an item and its expected lead time. As such, it is an estimate. A positive safety level indicates a willingness to invest money to be on the high side of that estimate when setting a reorder point. A negative safety level indicates a willingness to save some money and be on the low side of that estimate and use that money to invest in other items. No matter whether the safety level is positive or negative, the probability of a backorder will always exist because of the randomness of demand. The bottom line in setting safety-level quantities is balancing the risk against the investment to get the fewest backorders for the inventory investment.

Table 5-2 shows the relationship between the size of the safety-level quantity and various item attributes. (Because of their complexities, the mathematical relationships are not included in Table 5-2 .)

Since the safety level is designed to cover variance in lead-time demand, the fact that demand variance and lead time are the top two factors in determining the size of the safety-level quantity is not surprising. The inclusion of average requisition size among the other factors may be surprising at first.[13] The proposed model focuses on filling requisitions and not just units demanded during

[12] Some modelers have sought to develop a lead-time variance by expanding the sample size from all procurements for a single item to all procurements for a similar class or group of items and then applying that variance to each item in the group. Although the power rule that we suggested for extending demand variance over a lead time does not directly consider lead-time variance, it could be adapted to do so by varying the *t* factor.

[13] Average requisition size is determined by dividing historical units demanded by the number of requisitions for those units.

a lead time. Consequently, if all other factors are equal for two items, the one with the smaller average requisition size gets the larger safety-level quantity. Two considerations support that modeling bias toward average requisition size:

- More essential equipment items, which merit more protection, have higher prices and higher reliability and consequently would be requisitioned in smaller quantities

- Even a smaller safety-level quantity would provide a partial fill for an item with a larger requisition size.

Table 5-2.
Safety-Level Quantity (SLQ) and Item Attributes

Item attribute	Relationship to EOQ	
	Conversational	Logical
Demand variance (MAD	Directly proportional	If MAD ↑ then SLQ ↑↑↑ and if MAD ↓ then SLQ ↓↓↓
Lead time (L)	Directly proportional	If L ↑ then SLQ ↑↑ and if L ↓ then SLQ ↓↓
Average requisition size (Z)	Inversely proportional	If Z ↑ then SLQ ↓ and if Z ↓ then SLQ ↑
Unit price (C)	Inversely proportional	If C ↑ then SLQ ↓ and if C ↓ then SLQ ↑
Demand (D)	Inversely proportional	If D ↑ then SLQ ↓ and if D ↓ then SLQ ↑

Note: SLQ = safety-level quantity.

The influence of demand is secondary; that is, it affects the economic order quantity and it, in turn, affects the size of the safety level. The larger the order quantity, the smaller the safety-level quantity. If we procure 3 years of demand when we order, then we are ordering once every 3 years and at risk for backorders one time every 3 years. If we procure 6 months of demand when we order, then we are ordering twice a year and at risk for backorders six times every 3 years. Hence, the need for safety level increases as we reduce order quantities and vice versa.

Repair Quantity and Repair Point

For depot-level reparables, repair and not procurement is the primary source of supply. An item's repair-cycle quantity is the level of stock maintained to satisfy demand while unserviceable units go through the repair process. It is

computed as the product of the item's regeneration forecast and its forecasted time to repair (or repair cycle time).

Once a decision is made that an item will be managed as a reparable, questions arise concerning induction into maintenance. One approach is to induct every returned unserviceable unit into depot repair as soon as it arrives from the customer. This is a commercial practice that may not apply to the Coast Guard, which normally inducts batches of units into commercial and government maintenance facilities.

Repair Quantity for Government Repair

The approach for initiating batch induction mimics the approach to procuring stock. When a reparable item's inventory position (i.e., the sum of the serviceable and unserviceables units on hand and units due in) reaches or goes below its reorder point, a procurement order is initiated. Similarly, when a reparable item's serviceable inventory position reaches or goes below its repair point, a repair order is initiated. An item's repair point consists of its repair-cycle quantity, its safety level, and any special levels.

Since repair is the primary source of supply for a reparable item, the quantity to be repaired – at least those repaired at government facilities – should be computed in a similar manner. It should balance the cost of placing a repair order against the cost of repairing unserviceable assets. Therefore, to derive an economic repair quantity (ERQ), we borrow from the classical economic-order-quantity formula for items with level demand patterns to derive Equation 5-25:

$$ERQ = \sqrt{\frac{2A_r D_u}{Hc_r}} \qquad \text{[Eq. 5-25]}$$

where

A_r = cost to place a repair order

D_u = demand forecast minus demands with unserviceable item returns (i.e., carcass return rate)

H = holding rate

c_r = cost to repair one unit.

For nonlevel-demand patterns, we can again use part-period balancing with A_r and c_r to determine the optimal repair quantity. If unserviceable units are condemned as part of the repair process (and not condemned upon receipt from the fleet), then the actual induction quantity must be increased to account for condemnations. For example, if the condemnation rate is 10 percent and the ERQ is 60 units, then the induction quantity should be 66. In practice, the actual

induction quantity is always constrained by the number of unserviceable units on hand when the repair point is breached.

Repair Quantity for Commercial Repair

For commercially repaired Coast Guard items, we found that repair quantities often have been in the lot sizes negotiated by contracting officers as the "best buy" based on cost per unit. That approach could be strengthened by expanding the cost per unit to consider not only the cost that the Coast Guard is paying for unit repair but also the cost of placing the repair order. That latter cost is the sum of the administrative costs of making the contract and the distribution costs for issuing and transporting the unserviceable units and receiving the serviceable units (including transportation costs if not part of unit repair cost). Using the ERQ model as a starting position and then using the quantity discount EOQ models (with the different offering of unit repair costs, the cost to place a repair order, and the adjusted demand forecast), the Coast Guard could determine the true "best buy."

Repair Point

The repair point is analogous to the reorder point in that it signals when repair should be initiated. When the serviceable inventory position for an item, i.e., the level of stock on-hand and due-in from procurement and maintenance drops below its repair point, the unserviceable items should be directed to repair. The repair point is computed as the sum of an item's repair-cycle quantity, its safety level, and any special levels. That is,

Repair point = repair cycle quantity + safety level + special levels

or mathematically,

$$rr = R_q d_r + s + l_s \qquad \text{[Eq. 5-26]}$$

where

rr = repair point

R_q = repair-cycle time quantity

d_r = demand satisfied through repair (forecasted quarterly regenerations)

s = safety level

l_s = special level.

Repair-Cycle Time Quantity

As discussed in Chapter 3 on forecasting, repair-cycle time refers to the time between induction of an unserviceable unit into repair and when the serviceable unit is received and ready for issue. To preclude any stock outage, an unserviceable unit must be inducted one repair cycle time before the serviceable unit is actually needed. The repair-cycle-time quantity (RCTQ) is the number of items that will be requested during the period equal to the repair-cycle time.

As shown in the repair point equation, the computation of an item's RCTQ is analogous to the computation of its PLTQ; that is, it is computed by multiplying the item's regeneration rate by its repair cycle time. Likewise, for an item with a nonlevel forecasted regeneration pattern, the same nonlevel procedure discussed for PLTQ is used to compute RCTQ.

Safety Level for Reparables

For reparable items, demand can be segmented between that satisfied by repair (i.e., regenerations) and that satisfied through procurement (i.e., total demand less regenerations). Likewise, the variability of lead-time demand is divided between that of demand less regenerations over a procurement lead time and that of regenerations over a repair lead time. Since repair lead times are normally smaller than procurement lead times,[14] the variability of demand to be covered by the safety level is actually less than that of consumable items for which only procurement lead time is considered.

To develop a new MADLT (MADLTnew) to cover demand variability over both procurement and repair lead times, we start by computing a MAD for reparable regenerations (MADG) using the same formula as MAD except that we work with regenerations and not demand. We can develop a MAD for reparable demand not satisfied by regeneration (MADD) using the following equation from the statistics of combined variance:

$$MADD = \sqrt{\frac{MAD^2 - P_R MADG^2}{P_D}} \qquad \text{[Eq. 5-27]}$$

where

P_R = percent of demand satisfied by regeneration

P_D = percent of demand not satisfied by regeneration.

[14] Reparable items frequently are not off-the-shelf items and the production portion of their procurement lead times is not simply vendor delivery time but the combination of manufacturing time and delivery time. Barring funding, resource, or other constraints that could delay a repair or procurement action, the normal assumption is that the time to repair an item is less than the time to manufacture it.

Using the power rule formula (Equation 5-21), we convert MADD to MAD over a procurement lead time for reparable items and MADG to MAD over a repair lead time for reparable items. Using the statistics for a combined variance again, we have the following equation:

$$MADLT_{new} = \sqrt{(MADLT_{old})^2 - P_R\,(MADG)^2\left(\frac{L^{2\beta}-R^{2\beta}}{F^{2\beta}}\right)} \qquad \text{[Eq. 5-28]}$$

where

$MADLT_{new}$	=	new MADLT
$MADLT_{old}$	=	normal MADLT
P_R	=	percent of demand satisfied by regeneration
$MADG$	=	MAD for regenerations
L	=	procurement lead time in months
R	=	repair cycle time in months
F	=	forecast interval (i.e., 3).

With the new MADLT, we can once again use the safety-level model previously given.

CONCLUSIONS AND RECOMMENDATIONS

- Demand-supported stockage consists of the following basic inventory levels:
  - The requisitioning objective, or point to which one orders (it is the sum of the economic order quantity and the reorder point)
  - The economic order quantity, or the quantity one orders
  - The reorder point, or point at which one places an order (it is the sum of the safety-level quantity and the procurement-lead-time quantity)
  - The procurement-lead-time quantity, or the quantity that will be requested from the time a procurement action is initiated until stock is received
  - The safety-level quantity, or the quantity of stock one maintains on-hand to protect against stockouts while waiting for a procurement to arrive

- For reparables, a repair quantity, or the quantity that is inducted into repair when unserviceable item repair is directed

- For reparables, a repair point, or point at which repair is directed(it is the sum of the repair-cycle and safety-level quantities)

- For reparables, a repair-cycle quantity, or the quantity that will be requested from the time a repair action is initiated until the repaired units are returned.

- The computation of these inventory levels should be based on costs, responsiveness or support to the customer, and item attributes.

We recommend that the Coast Guard adopt the following models:

- For economic order quantities:

  - For a single item with level demand and a dollar value of demand outside limits for multiple costs, the classical EOQ model.

  - For a single item with level demand and a dollar value of demand within the limits for multiple costs, the multiple cost EOQ model.

  - For a single item with nonlevel demand, the part-period balancing algorithm with the look-ahead feature.

  - For a group of items with level demand, the group EOQ model.

  - For a group of items with nonlevel demand, the group part-period balancing algorithm.

- For lead-time quantities (procurement and repair):

  - For items with projected level demand, multiply average demand rate by period of lead time to arrive at the quantity.

  - For items with projected nonlevel demand, estimate the point in the future demand pattern at which the procurement or repair action will be initiated and tabulate the demand in the lead time that follows to arrive at the quantity.

- For safety level quantities, use a multi-item, response-time model with a maximum of three standard deviations of lead-time demand and a standard deviation of demand based on the mean absolute deviation.

- For reparable item repair quantities, use the same models as for the economic order quantities, except as variables use the cost to initiate a repair

action instead of the cost to order, use regenerations instead of demand, and use the cost of repair instead of the unit price of the item.

CHAPTER 6

Insurance and Low-Demand Stockage

INTRODUCTION

The Coast Guard must deal with a large percentage of low- or no-demand items (a low-demand item is an item with an average quarterly forecast less than one demand). If not managed as nonstocked items, these items pose special handling in forecasting and in setting levels.

Low-demand items often experience intervals of no demand interspersed with periods in which small quantities are required. Stockage is normally justified on the basis of their essentiality to either major system maintenance or mission readiness.

INSURANCE STOCKAGE

Insurance stockage is normally reserved for an item with no predicted demand but whose *failure without a spare would cripple mission performance.* For such items, we believe the inventory level should equal the minimum replacement unit. If a catastrophic failure does occur, then a spare is available to fill the need. If two or more are needed to reverse one catastrophic failure, the minimum replacement unit should reflect that number and they would be available to fill the need. Once a catastrophic failure occurs, an order for a replacement unit should commence with the issue of the unit on-hand.

LOW-DEMAND STOCKAGE

The computation of a stockage level for low-demand items has been a continuing challenge for inventory modelers. In practice, since none of the demand-based economic models apply, simple numeric levels are assigned to low-demand items. Demand over an extended period of time (e.g., 1 year) or the current level of assets is often used as the stockage level for a low-demand item. This approach does not consider the cost or the demand variability of the individual item and it is not directed at achieving a performance goal.

A Suggested Approach

An alternative approach to the use of numeric levels that addresses those problems is available. In developing inventory models, military and commercial

inventory researchers frequently use the Poisson probability distribution to represent demand frequency for low-demand items. If the number of demands is Poisson-distributed and each demand has a size of one,[1] then demand over a time interval is also Poisson-distributed. The mathematical tractability of the Poisson distribution yields solutions for modeling problems that seek to minimize costs while meeting either fill rate, response time, or operational availability goals.

Specifically, the probability of exactly n demands for a given average number of demands, λ, [p(n,λ)], is as follows:

$$p(n, \lambda) = \frac{\exp(-\lambda)\lambda^n}{n!} \qquad \text{[Eq. 6-1]}$$

where

exp = exponential function

λ = forecasted number of demands over lead time; item's daily demand rate multiplied by its lead time

n! = n factorial (i.e., if n = 4, 4! = 1 x 2 x 3 x 4 = 24

For a specific level of inventory, s, then the probability of all demands being issued immediately (or fill rate) is given by:

$$\textit{Probability of immediate issue} = \sum_{n=0}^{s-1} p(n, \lambda) \qquad \text{[Eq. 6-2]}$$

where

Σ = summation over index n, n = 0, 1, 2, ..., s-1 p(n,λ) and λ as above.

The expected number of backorders on hand is given by:

$$EBO(s) = \lambda - s + \sum_{n=0}^{s-1} [\,(s-n)p(n, \lambda)] \qquad \text{[Eq. 6-3]}$$

where

λ, s, Σ, and p(n,λ) as above.

In a normal processing line, the expected number in the line is equal to the expected number entering the line multiplied by the average time in the line.

[1] If the average requisition size for an item is greater than one, we can still apply the distribution if we consider it a requisition distribution versus a quantity distribution. Then, our stock level is in terms of requisitions and would need to be multiplied by the average requisition size to arrive at a quantity.

AD-A284 294 IMPROVED INVENTORY MODELS FOR THE UNITED STATES COAST GUARD REQUIREMENTS DETERMINATION PROCESS(U) LOGISTICS MANAGEMENT INST BETHESDA MD G L SLYMAN ET AL. OCT 93 2/2

UNCLASSIFIED LMI-CG-201RD6 XH-XD MDA903-90-C-0006 NL

END
FILMED
DTIC

Equivalently, the average time in the line is equal to the expected number in the line divided by the expected number entering the line. Applying those rules to requisition processing, the average time to fill a requisition (i.e., the mean system response time) is equal to the expected number of backorders on hand divided by the expected number of requisitions. (For example, if we average 20 backorders on hand and receive an average of 10 requisitions a day, our mean system response time is 20 divided by 10 or 2 days. Note that this assumes that requisitions that are filled immediately require zero time.)

We could use these expressions to formulate an inventory problem that reflects our objective. For example, we could formulate a multi-item inventory problem whose objective would be to meet a total-system-backorder on-hand goal at the lowest cost. It would be as follows:

$$\textit{Objective: } \min \sum_i C_i S_i \qquad \text{[Eq. 6-4]}$$

$$\textit{Constraint: } \sum_i \frac{EBO_i(S_i)}{Z_i} \leq \beta \qquad \text{[Eq. 6-5]}$$

where

- i = 1 to the number of low demand items
- C_i = cost of Item i
- S_i = stockage level or level of spares for Item i
- $EBO_i(S_i)$ = expected backorders on-hand for Item i for a given level of stock
- β = system backorder goal
- Z_i = average requisition size for Item i.

The system backorder goal is the sum of demands across all items times the target response time. (For example, if our total demand across all low-demand items is 3.24 demands per day and our response time goal is 30 days, our system backorder goal would be 97.) We can solve the problem using marginal analysis and arrive at our stock levels. To illustrate this approach, we prepared an example, which is shown in Appendix D.

Although this approach is not found in practice per se, it is the same approach found in readiness-based sparing models. We suggest it as an approach for managing low demand items at a Coast Guard supply center.

CONCLUSIONS AND *RECOMMENDATIONS*

- Setting levels for low-demand or no-demand items represents a special problem, since none of the traditional, demand-based economic models apply.

For insurance items or items with no forecasted demand, we recommend a minimal level of one replacement unit.

For items with low forecasted demand, we recommend the following:

- The use of the Poisson distribution to represent the expected demand pattern for an item.

- The use of a multi-item optimization model that seeks to provide an overall responsiveness goal at the lowest cost in inventory.

CHAPTER 7

Retention and Redistribution

INTRODUCTION

The need for retention models arises from three conditions:

- The phasing out of old equipment
- The planned obsolescence of an item being replaced with a newer one
- The declining demand caused by a change in operating procedures generating fewer failures or a change in maintenance procedures resulting in a less urgent requirement to replace the item.

Whether because of declining demand, item obsolescence/phasing out, or equipment phasing out, on-hand and on-order stock can become excess to normal requirements. When that occurs, the question is, "How much of the stock should be retained and how much disposed of?" The answer to that question is determined by the retention limit.

Where retention models deal with excess supply center stocks, redistribution models handle excess stocks in units supported by the supply center. Besides reports on excess supply center stocks, supply center inventory managers receive excess asset reports submitted by unit supply departments and have to make decisions on the disposition of these redistributable assets. Managers may retain some or all such assets in place or they may direct that the assets be returned to supply center storage sites or be sent to disposal. The role of the redistribution model is to aid them in making that decision.

RETENTION LEVEL

The phase-out portion of an equipment's life cycle presents significant challenges to the requirements determination process and the inventory manager. Replacement parts requirements must be balanced against declining equipment density as overhaul programs wind down and cannibalization of parts increases. Parts procurement to keep aging equipment operating represents dollars that may be wasted. On the other hand, limiting parts procurements can cause more expensive emergency purchases or locally applied modifications to keep the equipment's operating unit replaced. Some inventory managers consider the phase-out period for old equipment even more intense than the phase-in period for new equipment.

In the case of equipment being replaced, the retention limit for an item associated with the equipment is the expected demand for the item until phaseout is complete. Since the dates for item phaseout or equipment phaseout may change and demand forecasting is uncertain, some hedging is acceptable in setting retention limits for items associated with phaseout. Also, since demand will decline as an equipment is being phased out, the demand pattern for the item should reflect declining demand as it approaches the phase-out date.

Similarly, in the case of an item being replaced, the retention limit for the item is the expected demand for that item until it is replaced with the newer item. Assets above the limit can either be sent to disposal or issued in place of the newer item whenever that issuance is acceptable.

For items associated with equipment that is not being phased out and are not being phased out themselves, the retention decision is concerned with assets that are excess to supply center requirements. In this case, the retention limit for an item sets the point below which stock for the item should be retained and above which stock should be offered for disposal. The computation of this retention limit should be based on economic and/or noneconomic criteria. Normally, for ease of use, it is expressed in terms of years of stock.

Economic Retention

An economic retention model (ERM) balances the cost of holding a marginal unit of stock against the cost of not holding the unit to arrive at the maximum number of years of stock to hold. The economic retention limit for an item is defined as the last year that the cost of retention is less than or equal to the cost of disposal. Some Government agencies use economic limits expressed in years of stock and those limits are the same for the entire item population. Although the limits may be based on an ERM, a better approach would be to apply an ERM on an individual-item basis, similar to the way the EOQ model is applied.

In an ERM, the cost of retaining a unit of stock for j years is the combined cost of storage and investment weighted by the probability that the unit will be stocked (i.e., not used to satisfy demand). The cost of not retaining the unit is the cost of disposal less the net return from disposal and the cost of reprocurement if needed to satisfy a backordered demand. The ERM model takes the form of a break-even analysis represented by the following equation:

$$\sum_j [d_j SCO_j] = \sum_j [d_j P_j (A_u + (1+E_j)C] - (NC - C_d) \qquad \text{[Eq. 7-1]}$$

where

j	=	Years 1 through infinity (25 years is sufficient)
d_j	=	discount rate for year j, which is based on the investment or cost-of-capital portion of the holding rate
S	=	storage and loss portion of holding rate
C	=	unit price of item
O_j	=	probability of stock on-hand in Year j, which is $1 - F(x,j)$ where $F(x,j)$ is the probability that the cumulative demand for j years is equal to or greater than x units (x is expected demand in j years)
E_j	=	price escalation factor for buying item in Year j (1 = no escalation above normal inflation)
A_u	=	unit cost to order, which is cost to order (A) divided by EOQ
P_j	=	probability of a procurement in Year j, which is $F(x,j) - F(x, j - 1)$
N	=	net disposal rate
C_d	=	cost of a disposal.

The retention limit is the x (i.e., years of stock) when the left-hand side of the equation is approximately equal to the right-hand side of the equation.

For example, suppose the variables associated with our ERM had the following values for an item:

Cost of item (C)	=	$100
Annual demand (D)	=	100 units (used to compute EOQ)
Cost to order (A)	=	$500
Cost to hold (H)	=	16 percent (used to compute EOQ)
Investment rate	=	8 percent [$d_j = 1/(1+.08)_j - 1$]
Obsolescence rate	=	4 percent[1]

[1] For the sake of simplicity, we did not include obsolescence in our example. If we had, we would have adjusted the F(x,n) probabilities to include the probability that demand for an item may not materialize because of obsolescence.

Storage & loss rate (S) = 4 percent

EOQ = 79 units

Unit cost to order = $6.33

Net disposal rate (N) = 10 percent

Unit disposal cost (C_d) = $10

Inflation rate = 4 percent per year [$E_1 = 1$ and $E_j = (1.04)(E_j - 1)$].

and a simple probability of cumulative demand[2] of 0.00, 0.25, 0.50, 0.75, and 1.00, etc., as follows:

x in years	F(x,1)	F(x,2)	F(x,3)	F(x,4)	F(x,5)	F(x,6)	F(x,7)
1	0.50	0.75	1.00	1.00	1.00	1.00	1.00
2	0.25	0.50	0.75	1.00	1.00	1.00	1.00
3	0.00	0.25	0.50	0.75	1.00	1.00	1.00
4	0.00	0.00	0.25	0.50	0.75	1.00	1.00

The interpretation of the $F(x,n)$ probability is that we have a 0 percent probability that n years of demand will occur by Year n-2, a 25 percent probability that n Years of demand by Year n-1, a 50 percent probability by Year n, a 75 percent probability by Year n+1, and a 100 percent probability by Year n+2. Using the $F(x,n)$'s, we can compute the O_j's and P_j's for Unit x where x is a number of stocks (i.e., in our example, 100, 200, 300, 400, 500, etc.).

For the values of A and H (and its subelements), we used Supply Center, Curtis Bay's *Cost to Order and Cost to Hold Study*[3]. Using all of our values, we computed the following costs for the left-hand side (LHS) and right-hand side (RHS) of the ERM for increasing x years of stock. The computation for x equal 3 is shown in Table 7-1. Table 7-2 provides an example summary of the ERM for x from 1 to 20.

Based on these results, the economic retention limit for the item in our example is 19 years of stocks (the year in which the RHS and LHS are closest).

---

[2] This probability distribution is only for purposes of our example. For actual implementation, the normal probability distribution or an empirical distribution would be the likely candidate for most items and the Poisson distribution for slow moving items.

[3] *Cost to Order and Cost to Hold Study*, Fleet Support Operations Branch, USCG Supply Center, Curtis Bay, 1 June 1992.

Table 7-1.
Samples of ERM Computation

Year (n)	LHS					RHS			
	d_n	F(3,n)	O_n	SC	d_nSCO_n	P_n	$\acute{E}_n = 1 + E_n$	$C_d - NC$	$d_nP_n (A_u+\acute{E}_nC)$
1	1.000	0.00	1.00	$4.00	$4.00	0.00	1.00	$(9.90)	$0.00
2	0.926	0.25	0.75	$4.00	$2.78	0.25	1.04	$0.00	$25.54
3	0.857	0.50	0.50	$4.00	$1.71	0.25	1.08	$0.00	$24.54
4	0.794	0.75	0.25	$4.00	$0.80	0.25	1.12	$0.00	$23.68
5	0.735	1.00	0.00	$4.00	$0.00	0.25	1.17	$0.00	$22.66
6	0.681	1.00	0.00	$4.00	$0.00	0.00	1.22	$0.00	$0.00
Value					$9.29				$86.42

Table 7-2.
ERM Example

Year (n)	LHS	RHS	Year (n)	LHS	RHS
1	$2.93	$93.34	11	$29.84	$60.27
2	$5.71	$90.34	12	$31.63	$57.57
3	$9.29	$86.42	13	$33.29	$54.97
4	$12.60	$82.66	14	$34.82	$52.48
5	$15.67	$79.05	15	$36.24	$50.09
6	$18.50	$75.59	16	$37.56	$47.79
7	$21.13	$72.27	17	$38.78	$45.58
8	$23.57	$69.08	18	$39.90	$43.46
9	$25.85	$66.02	19	$40.95	$41.42
10	$27.91	$63.09	20	$41.92	$39.46

Noneconomic Retention

Noneconomic factors, such as contingency planning or the potential future loss of a source of supply, can also affect the retention decision. In the case of contingency planning, the quantity to be retained comes from outside the

inventory system and is merely an input to the system. In the case of the potential or actual loss of a source of supply, a life-of-type level can be determined by extending the demand forecast for the item through its expected life in the system. A reduction of demand to zero should be part of the end of the item's lifetime demand pattern.

Another potential reason for overriding the economic retention limit is the potential for significant changes to the future demand pattern of the item. If an item manager knows that the long-term forecasted demand pattern for the item is incorrect because of increasing maintenance or a decreased number of end items to be supported or some combination of both conditions, the manager may extend or reduce the retention limit accordingly.

Any overriding by the item manager is always an important consideration. Circumstances outside the scope of an economic model will occasionally arise. To deal with them, the item managers must be permitted to retain assets above the limit set by an economic model.

Redistributable Asset Disposition Decision Models

When a unit's supply department reports that it has excess stock of an item, the response of the Supply Center can be to return the stock to the supply center storage locations, to retain the stock in place, or to dispose of the stock. The exact disposition that the supply center transmits to the supply department depends on the total asset position of the item; that is, the assets in the supply center storage locations plus the sum of reported redistributable assets:

- If the total asset position is above the item's retention level, the return disposition is not a viable option and the disposition decision devolves to either retain assets in place or dispose of redistributable assets (i.e., dispose of supply center assets or dispose of site assets).
- If an item's total asset position is below its retention level, disposal is not a consideration and inventory managers must decide to retain the assets in place or return redistributable assets.
- If an item's total asset position is below its retention levels but the supply department indicates that retaining assets in place is not a viable option because of limited storage capacity, the disposition decision is reduced to returning or disposing of the redistributable assets.

For each of these three situations, we will establish decision models that are suitable for use by inventory managers.

Retain-In-Place/Disposal Model

The retain-in-place/disposal model should consider the following factors in the given sequence.

- *Diminishing manufacturing sources* – The first factor to evaluate is based on a support consideration, namely, the sources of supply for active items in the Coast Guard inventory are diminishing. Redistributable assets should be retained for items that have been approved for life-of-type buys.

- *Economic factors* – On the basis of economic considerations, the Coast Guard must determine whether some or all redistributable assets should be retained or disposed of. Where an item's total asset position is such that a portion of stock is excess, the Coast Guard must decide whether the level of excess for disposal should be drawn from assets at the supply center storage sites or from redistributable assets at unit sites. The cost tradeoff is between disposal at the supply center sites and/or disposal at one or more unit sites. The number of disposal actions, the positioning of assets relative to future demand, and the potential future use of a redistributable asset at a unit site should all be key elements in making this tradeoff.

- *Noneconomic factors* – A normal assumption in logistics modeling is that demand is known. Indications that demand over the retention period will be different than the item's forecast should alert inventory managers to determine whether increasing or decreasing demand activity justifies an override of a disposition from the economic model with justification.

Although models should be flexible, they will not be able to consider all possible conditions properly. Inventory managers should have sufficient authority to override a disposal decision from an economic model.

Retain-In-Place/Return Model

The retain-in-place/return model should consider the following factors in the given sequence.

- *Economic factors* - Again, it is necessary to determine whether some or all of the assets should be retained in place or returned to a wholesale storage site based on economic considerations. The cost tradeoff is between returning the assets and issuing them from a supply center storage site and retaining the assets in place and issuing them from a unit site. The cost of returning an asset, the positioning of assets relative to future demand, and the potential future use of a redistributable asset at a unit site should all be key elements in making that tradeoff. In general, if the cost of storing and issuing an asset from a supply center storage location and from a unit site are the same, then the only cost is that incurred by the return; in that case, the decision should always be to retain the asset in place.

- *Item managers' override* — Although models should be flexible, they are not able to consider all possible conditions properly. For example, a model may not be able to distinguish between redistributable assets at an accessible site (i.e., a site with several transportation channels to other sites) and redistributable assets at a remote site (i.e., a site outside of normal transportation channels). Returning assets to a supply center storage site from a remote site may be preferable to a model disposition of retain in place. Inventory managers should have sufficient authority to override a disposition from an economic model with justification.

Return/Disposal Model

The return/disposal model should consider the following factors in the given sequence.

- *Diminishing manufacturing sources* — The first factor to evaluate is based on a support consideration, namely the sources of supply for active items in the Coast Guard inventory are diminishing. Redistributable assets should be returned for these items.

- *Economic factors* — On the basis of economic considerations, the Coast Guard must determine whether some or all of the assets should be returned. If a site cannot retain an item's redistributable assets in place, the cost tradeoff is between returning and issuing those assets from a supply center site and the disposal of those assets at the unit site. That tradeoff is identical to that made by the ERM except that the left-hand side of the equation has the additional cost of the return. Hence, the procedure is to determine a new optimum, j', that includes the cost of the return. If the supply center asset position is above j' but below j (the retention level), then the decision would be to dispose and not return the assets. If the asset position is below j', the decision would be to return the stock.

- *Noneconomic factors* — The same noneconomic factors apply for the return/disposal model as apply for the retain-in-place/disposal model.

RECOMMENDATIONS

For retention levels, we recommend the following:

- The retention limit for an item associated with equipment being replaced should be the expected demand for the item until the equipment is phased out. (Some hedging is acceptable.)

- The retention limit for an item being replaced should be the expected demand for that item until it is replaced with the newer item.

- The retention limit for items not associated with equipment being phased out or not being phased out themselves should be set at the level at which the cost of holding the stock is equal to the cost of not holding it.

We recommend that unit reported excesses be redistributed as follows:

- If the total asset position for the item is above its retention level, then a retain-in-place/disposal model should be used to determine the item's disposition.

- If the total asset position is below the item's retention level, then a retain-in-place/return model should be used to determine the item's disposition.

- If the total asset position is below the retention level but the supply department indicates that retaining the asset in place is not a viable option, then a return/disposal model should be used to determine the item's disposition.

- All three models should consider both economic and noneconomic factors and provide for item manager discretion.

APPENDIX A

Readiness-Based Sparing Models

BACKGROUND

In the past 20 years, two major developments have occurred in inventory management. The first is the emergence of the just-in-time (JIT) concept and the second is readiness-based sparing (RBS). In the Department of Defense (DoD), the Military Services are using RBS in provisioning spares for major weapon systems and, to some extent, in replenishing those spares. For example, the Air Force is using the Aircraft Availability Model to manage its wholesale and retail inventories of depot-level reparable aviation items, and NASA is exploring its use in supporting spacecraft. In the private sector, it is being promoted as a means of supporting power plants, production equipment in manufacturing plants, and telecommunications systems.

WHAT IS IT?

Readiness-based sparing is the process by which inventories are sized to support a major system or principal end item at levels that meet the customer's availability goal for operating that system or end item. Appropriately, RBS is sometimes referred to as sparing to availability. An RBS model is an inventory requirements model that sets levels that either minimize costs while meeting a system or end-item operational availability (A_o) goal or maximize the system or end item's A_o within a dollar limit.

As items (parts) fail on a weapon system or end item, the customer places demands on the supply system for replacements. The weapon system or end item is considered operational or ready when it can perform its assigned job(s) and not operational when an item failure prevents it from doing so. Therefore, the time that the supply system takes to fill customer demands affects the readiness of the equipment the shorter the time, the less time the system or end item is not operational.

The level of inventory for an item affects the time the supply system needs to respond to a customer's demand for that item in two ways. First, it governs the number of demands that are filled immediately (i.e., the item's fill rate or supply availability). Second, it limits the time that the customer must wait for demands that are not filled immediately (i.e., the item's backorder time). The combination of these two measures is an item's average customer wait time

(ACWT) or mean system response time (MSRT). Its calculation is shown in Equation A-1:

$$ACWT = (f)(k) + (1 - f)(b_t) \qquad \text{[Eq. A-1]}$$

where

f = fill rate

k = time to make an immediate issue

b_t = average time on backorder.

(A comparison of this equation and Equation 2-2 would reveal that MSRT and ACWT are one and the same measure of supply support.) If the time to make an immediate issue is assumed to be zero and since 1- f is the backorder rate, which is also given by the number of backorders established over total orders, we have Equation A-2:

$$ACWT = (1 - f)\,(b_t) = \left(\frac{EBE_i}{d_i}\right) b_t = \frac{EBO}{d_i} \qquad \text{[Eq. A-2]}$$

where

EBE_i = expected backorders established

d_i = total orders or demand for item

EBO_i = expected backorders on-hand.

Like a response-time model, an RBS model targets on an item EBO. However, unlike a response-time model, an RBS model seeks the best combination of item EBOs in terms of the availability of the system or end item being supported.

A number of alternative mathematical expressions exist for defining a system's or end item's operational availability. One is based on the premise that each backorder represents a "hole" in a system or end item and causes it to be considered "not ready". If the failure of an item is independent of failures for other items, we know from probability theory that the probability of system or end item having no failures (i.e., ready) is the product ($\prod$) of the probabilitiesthat

each item is ready. Accordingly, the operational availability expression is that shown as Equation A-3:

$$A_o = \prod_i \left(1 - \frac{EBO_i}{NS}\right)^{QPS_i} \qquad \text{[Eq. A-3]}$$

where

A_o = system or end item operational availability

i = item index from one to the number of items in system

EBO_i = expected backorders on hand for item i

NS = number of systems

QPS_i = quantity of item i per system.

How Does an RBS Model Work?

Marginal analysis is the optimization approach used by most RBS models. The marginal analysis approach is follows:

- For each item, calculate a cost-benefit ratio for adding spares to the item's inventory level.
- Rank the items by their ratios.

$$Ratio = \frac{A_o \textit{ improvement if spare is added}}{\textit{Cost of spare}} \qquad \text{[Eq. A-4]}$$

- Select the spare with the most improvement per dollar.
- Recalculate the cost-benefit ratio for the next spare for the item selected and rerank the item according to that ratio.
- Repeat the previous two steps until the A_o goal is reached or until the dollar limit is reached.

The marginal analysis approach does not yield the typical closed-form solution (i.e., an equation defining the solution) that is found in other optimization techniques such as the Method of Lagrange Multipliers. However, as demonstrated by T.J. O'Malley in a Logistics Management Institute (LMI) report, the list of spares from this approach is optimal.[1]

[1] *The Aircraft Availability Model: Conceptual Framework and Mathematics*, O'Malley, T.J., LMI Report No. AF201, June 1983.

Other aspects of RBS modeling are multi-echelon and multi-indenture capabilities. If an RBS model has a multi-echelon capability, its selection process considers the best location for that spare among all the echelons of supply supporting the system or end item. That approach entails developing separate ratios for the spare at each echelon and picking the best ratio. Since operational availability support starts at the first level of supply supporting a system or end item, a single-echelon RBS model would be a "retail" model. By definition, a "wholesale" single-echelon RBS model is not possible. "True" multi-echelon RBS modeling has only been performed for reparable items. For consumable items, only single-echelon or pseudo-multi-echelon (upper-echelon support set by tradeoff analysis with lower-echelon support) RBS modeling has occurred.

If an RBS model has a multi-indenture capability, its selection process considers the position of the item in the system or end item's configuration. A first-indentured item goes directly on the end item, while a second-indentured item goes on a first-indentured item. If an item is at the top of the hierarchy (i.e, is a first-indentured item), its impact on operational availability is more direct and consequently greater than if it is lower in the hierarchy (i.e., second-, third-, etc., indentured items). For example, a spare for a second-indentured item will reduce the repair time for its associated first-indentured item and that will, in turn, reduce the time the system or end item is not available when the first-indentured item fails. On the other hand, a spare for the first-indentured item will directly reduce that time.

The multi-indenture capability is incorporated by considering hierarchy in the computation of a spare's operational availability improvement. Once the marginal analysis algorithm selects a spare in a hierarchical chain, the ratios for all items in the chain must be adjusted to reflect their new operational availability contribution.

What Is Required for RBS Modeling?

The following data are required to run a multi-echelon, multi-indenture RBS model:

- A complete configuration of the system or end item listing the indenture and essentially of each item, starting with first-indentured items down to the last indenture for which you want to calculate RBS
- For each item, its failure rate, its unit price, and its recoverability code (i.e., consumable or reparable designation)
- For each reparable item, the level of repair and for each level of repair, the repair times and percentage repaired, the percentage sent to the next higher level of repair, and the percentage condemned.
- The transit times between echelons of supply

- The overall operational availability goal and dollar limitation, if one exists, for the major system or end item (note that site-specific goal setting is still in development).

RBS modeling also requires the following:

- Specification of a probability distribution for failure rates and repair times (normally assumed to be exponential so that the probability density function for the number of items in the pipeline is Poisson)

- An operational availability expression such as the product rule previously defined

- A cutoff mechanism to ensure that the tail of the demand distribution for low-cost items is not driving the spares solution (e.g., stop sparing when item's fill rate is 99.9 percent).

An Example

A (fictitious) equipment comprises four items with the attributes shown in Table A-1. If we assume a Poisson distribution for failures, we compute the spares availability data shown in Table A-2. Given the spares data in the marginal analysis, the algorithm would select a spare for Item 3 since it has the largest improvement for the dollar. The algorithm would then see that the operational availability or dollar target was reached. If not, it would compute Item 3's potential availability with two spares and use it to compute a new potential equipment operational availability and new improvement-to-cost ratio if the next spare were for Item 3. It would again compare the improvement-to-cost ratios against the four items and select the item with the largest improvement for the dollar. It would repeat this process until the equipment operational availability or budget dollar target is reached.

Table A-1.
Item Attributes

Item attribute	Item 1	Item 2	Item 3	Item 4
Cost per unit quantity	$1,000	$500	$1,000	$500
Failures per year (100 equipment)	50	10	50	10
Resupply time (to organizational unit)	20 days	20 days	20 days	20 days

Table A-2.
Spares Availability Data

	Item 1	Item 2	Item 3	Item 4
Item avail. with 0 spares	0.9726	0.9945	0.9726	0.9948
Equipment A_o	0.9356	0.9356	0.9356	0.9356
Item Avail. with 1 spare	0.9820	0.9987	0.9987	0.9987
Equipment A_o	0.9446	0.9446	0.9446	0.9446
Improvement/cost	0.000008	0.000008	0.000008	0.000008

APPENDIX B

Forecasting Models

This appendix describes the variety of stochastic models available for forecasting. All are time-series models, that is, they all focus entirely on the historical demand pattern to generate a forecast. The models are presented within the following general category or class of models:

- Simple time-series models
- Smoothing models
- Linear-trend models
- Nonlinear-trend models
- Decomposition methods
- Seasonal smoothing models
- Box-Jenkins method
- Combined forecasts.

To help understand how the models perform with different data patterns, we constructed the example shown in Table B-1 using five demand patterns for which we have eight quarters of demand history plus the next quarter. (The item in Table B-1 with the random pattern is an actual Coast Guard item with its eight quarters of historical demand plus, for the next period, the rounded average of the eight quarters.) We use the next-period quantity as a basis for comparing forecasts to actual demand to help judge the error produced by the different models.

Table B-1.
Demand History (Example)

Pattern	8	7	6	5	4	3	2	1	Next period
Level	10	10	10	10	10	10	10	10	10
Trend	10	15	20	25	30	35	40	45	50
Seasonal	10	20	30	20	10	20	30	20	10
Cyclical (5-period)	30	10	20	25	15	30	10	20	25
Random	80	18	135	2	8	4	24	14	36

Simple Time-Series Models

Description

Some time-series models are quite complex; others are simple and easy to understand, and those are the ones we consider here. Simple time-series models are sometimes characterized as being naive models because they use the most basic assumptions on how future values of the time series can be predicted with past values. The four models within this category are described in Table B-2.

Example

If we use the four simple models to forecast the next quarter for our five data patterns, some of the computations would be as follows:

- *Basic model with level pattern* – Since the actual demand for Quarter 1 (the last period) was 10, the forecast for the next quarter is 10.
- *Basic seasonal model with seasonal pattern* – Using a four-quarter version of the model, we look at the actual demand for Quarter 4 (a year ago), which was 10; thus, the forecast for the next quarter is 10.
- *Levels model with random pattern* –The sum of the eight past quarters is 285. The forecast for the next quarter is 285 divided by 8 or 35.6.
- *Change model with trend pattern* – Since the seven changes between each quarter and the succeeding quarter for eight quarters are all +5, the average

Table B-2.
Simple Models

Model	Formulation of forecast
Basic model	Forecast = actual value for last period
Basic seasonal model	Forecast for 1st period = actual value for last period
	Forecast for 2nd period = actual value for 2nd period
	Forecast for 3rd period = actual value for 3rd period, etc., until forecast for last period = actual value for last period
	"Period" is the forecast period (e.g., a quarter) and the number of periods depends on the seasonality (e.g., 4 quarters)
Levels model	Forecast for next period = total sum of all past periods divided by the number of past periods
Change models	Forecast for next period = actual value for last period plus average change
Versions are:	where
Average change	average change = average of changes
Percent change	average change = average percentage change times last value
Weighted change	average change = weighted average of changes

change is +5. Using the average change model, the forecast is 45 (the last period) plus 5, or 50.

The results of all of the computations are summarized in Table B-3. For that example and all examples in this appendix, the forecasts are rounded to the nearest tenth.

Table B-3.
Simple Model Sample Computations

Pattern	Model and forecast				Actual value
	Basic	Basic seasonal	Levels	Average change	
Level	10	10	10	10	10
Trend	45	30	27.5	50	50
Seasonal	20	10	20	21.4	10
Cyclical	20	15	20	18.6	25
Random	14	8	35.6	4.6	36

Smoothing Models

Description

The smoothing models assume that the time series consists of a level pattern plus fluctuations caused by randomness. Models in this category attempt to smooth out the fluctuations by smoothing or averaging them. Like the simple models, these models are easy to use and relatively easy to understand. In fact, they use fewer data than the simple change models. The two main models in this category are described in Table B-4.

Table B-4.
Smoothing Models

Model	Formulation of forecast
Moving average	Forecast for next period = sum of actual values for some number of past periods divided by the number of past periods
Single exponential smoothing	Forecast for next period = $(\alpha) \bullet$ (last period) + $(1 - \alpha) \bullet$ (last forecast) where α (alpha) = smoothing constant

Example

If we used the two smoothing models to forecast the next quarter for our five data patterns, some of the computations would be:

- *Moving average model with trend pattern* –

  - For a four-quarter moving average, the forecast is the sum of the last four quarters (i.e., 30 + 35 + 40 + 45 = 150) divided by 4, or is 37.5.

  - For an eight-quarter moving average, the forecast is the sum of the last eight quarters (i.e., 10 + 15 + 20 + 25 + 30 + 35 + 40 + 45 = 220) divided by 8, or 27.5.

- *Single exponential smoothing model with seasonal pattern* –

  - Using $\alpha = 0.2$, we have a smoothed forecast of 18.9 at the end of seven quarters. The forecast is a combination of two quantities: 0.2 times the demand for the last period (20) and 0.8 times the previous forecast (18.9), which results in a forecast of 19.1.

  - Using $\alpha = 0.1$, we have a smoothed forecast of 15.5 at the end of seven quarters. The forecast again is a combination of two quantities: 0.1 times 20 (the demand for the last period) and 0.9 (1 – 0.1) times 15.5 (the previous forecast), which results in a forecast of 16.0.

The results of all of the computations are summarized in the following Table B-5. As would be expected, these models do well with patterns that approximate level but do not do as well as the simple models when it comes to trend patterns.

Table B-5.
Smoothing Models Sample Computations

Pattern	Model and forecast				Actual value
	Moving average (4 qtrs.)	Moving average (8 qtrs.)	Single smooth 0.2)	Single smooth (0.1)	
Level	10	10	10	10	10
Trend	37.5	27.5	29.2	21.5	50
Seasonal	20	20	19.1	16	10
Cyclical	18.8	20	21.2	24.1	25
Random	12.5	35.6	34.7	51.8	36

Linear Trend Models

Description

The linear trend models assume that the time-series consists of a upward or downward trend pattern plus fluctuations from randomness. The four models in this category, which are more complex than the simple or smoothing models, are given in Table B-6.

Table B-6.
Linear Trend Models

Model	Formulation of forecast
Linear regression	Forecast = $a\,T + b$ where T = time code for next period. For example, quarter 1 2 3 4 5 6 7 8 next code −7 −5 −3 −1 1 3 5 7 9 and $a = \frac{\sum_i d_i T_i}{\sum_i T_i^2}$ $b = \frac{\sum_i d_i}{n}$ where T_i = quarter time code d_i = demand for quarter n = number of quarters
Linear moving average	Forecast = $L + T_t$ where L = level term $= 2\,M_1 - M_2$ T_t = trend term $= [2/(N-1)]\,[M_1 - M_2]$ where M_1 = moving average of demand M_2 = moving average of M_1 N = number of periods in average

Table B-6.
Linear Trend Models Continued)

Model	Formulation of forecast
Double exponential smoothing	Forecast = 2 $S_1 - S_2 + T$ where $S_1 = (1 - \alpha)$ (previous S_1) + (α) (last period) $S_2 = (1 - \alpha)$ (previous S_2) + (α) (S_1) T = trend term $= [\alpha/(1 - \alpha)][S_1 - S_2]$ where S_1 = single smoothed average S_2 = double smoothed average T = trend term α = smoothing constant
Holt's two-parameter exponential smoothing	Forecast = $L_t + T_t$ where L_t = level term $= \alpha$ (last period) + $(1 - \alpha)$ $(L_{t-1} + T_{t-1})$ T_t = trend term $= \beta(L_t - L_{t-1}) + (1 - \beta)\ T_{t-1}$ where α = *1st smoothing constant* β = 2nd smoothing constant

Example

If we use the four trend models to forecast the next quarter for our five data patterns, some of the computations would be as follows:

- *Linear regression model with trend pattern* — The sum of the demands for the eight quarters is 220. The sum of the demands multiplied by their respective time codes is 420 [(10)(–7) + (15)(–5) + (20)(–3) + (25)(–1) + (30)(1) + (35)(3) + (40)(5) + (45)(7)]. The sum of the square of the time codes is 168. Therefore "a" is 2.5 (420/168) and "b" is 27.5 (220/8). With T equal to 9, the forecast is 2.5 multiplied by 9 plus 27.5, which results in a forecast of 50.0.

- *Linear moving average with trend pattern* — Using a four-quarter moving average, we have a first moving average of 37.5 and a second moving average of 25 (average of 1st averages of 17.5, 22.5, 27.5, and 32.5) at the end of eight quarters. The trend term is 2/3 [2/(4 – 1)] multiplied by 12.5 (37.5 – 25), which yields a term of 8.3. The forecast is 75 [2 times 37.5 (1st average)] minus 25 (2nd average) plus 8.3 (trend), which results in a forecast of 58.3.

- *Double exponential smoothing with trend pattern* – Using α = 0.2, we have a single smoothed average of 25.2 and a double smoothed average of 14.7 at the end of seven quarters. The new single smoothed average is 0.2 multiplied by 45 (the demand for the last period) plus 0.8 (1 - 0.2) multiplied by 25.2 (the previous average), which results in a new average of 29.2. The new double smoothed average is 0.2 multiplied by 29.2 plus 0.8 multiplied by 14.7, which results in a new average of 16.8. The trend term in the forecast is 0.25 (0.2/0.8) multiplied by 12.4 (29.2 - 16.8), which yields a term of 3.1. The forecast is 2 multiplied by 29.2 (the new single average) minus 16.8 (the new double average) plus the trend term, which results in a forecast of 44.7.

- *Holt's exponential smoothing with trend pattern* – Using α = 0.9 and β = 0.9, we had a linear term of 40.0 and a trend term of 5.0 at the end of seven quarters. The new linear term is 40.5 [0.9 (α) multiplied by 45 (last period)] plus 4.5 [0.1 (1 - α) multiplied by 45.0 (sum of previous linear and trend terms)] which yields a new term of 45. The new trend term is 4.5 [0.9 (β) multiplied by 5 (45 - 40)] plus 0.5 [0.1 (1 - β) multiplied by 5.0 (the previous term)], which results in a new term of 5.0. The forecast is 50 (45 plus 5).

Table B-7 summarizes the results of all computations. As would be expected, these models do well with trend patterns.

Table B-7.
Linear Trend Models Sample Computations

Pattern	Model and forecast					Actual value
	Linear regression	Linear moving average	Double smooth (0.2)	Double smooth (0.1)	Holt smooth (0.9, 0.9)	
Level	10	10	10	10	10	10
Trend	50	58	44.7	31.9	50	50
Seasonal	24.3	20	24.7	20.8	16.5	10
Cyclical	17.3	16.1	16	19.5	21.4	25
Random	0*	0*	2.3	27.5	61	36

Note: * = negative set to zero.

NONLINEAR TREND MODELS

Description

The linear trend models attempt to fit the data to a straight line that is a graph of a linear trend. The nonlinear trend models attempt to fit the data to other curves which are not linear trends. If *F* represents the forecast, *t* the time

period (1, 2, 3, etc.), and a, b, and c parameters, we have the following curves and their associated forecasting models:

- Linear trend $F = at + b$
- Inverse linear $F = a/t + b$
- Exponential curve $F = b\ exp(at)$
- Compound growth $F = b\ a^t + c$
- Modified exponential $F = b\ t^a + c$
- Logistic $F = 1/(b\ a^t + c)$

A number of other curves that are combinations of the above exist.

Model

To illustrate how this category of models, we selected one model to work with – the exponential curve. Starting with $F = b\ exp(at)$, we take the natural logarithm (ln) of both sides and get:

$$\ln(F) = \ln[b\ exp(at)] = \ln(b) + \ln[exp(at)] = \ln(b) + at$$

If we let $b' = \ln(b)$, we have $\ln(F) = at + b'$. This is the form of the linear regression model. Working with the logarithm of the demand for the past periods and using the same equations we did for that model, we can solve for a and b'.

Example

We can use the exponential curve models to forecast the next quarter for the trend pattern. First, the sum of the logarithms of the demands for the eight quarters is 25.677. The sum of the logarithms of the demands multiplied by their respective time codes is 17.294 [(2.302)(–7) + (2.708)(–5) + (2.995)(–3) + (3.218)(–1) + (3.401)(1) + (3.555)(3) + (3.688)(5) + (3.806)(7)]. The sum of the square of the time codes is 168. Therefore, a is 0.103 (17.294/168) and b' is 3.21 (25.677/8). With T equal to 9, the forecast is the exponential of 0.103 multiplied by 9 plus 3.21, which results in a forecast of 62.6. The results for all patterns are summarized in Table B-8.

Table B-8.
Nonlinear Trend Model Sample Computation

Pattern	Model and forecast (exponential curve)	Actual value
Level	10	10
Trend	62.6	50
Seasonal	24.3	10
Cyclical	16.5	25
Random	5.5	26

Decomposition Method

Description

The underlying assumption in the decomposition method is that the data pattern comprises four components – a trend component (T), a seasonal component (S), a cyclical component (C), and a random component (I). The decomposition method attempts to isolate these components in the historical time series and then recombine them into a forecast for the future. The formulations of the forecasts by two common decomposition models are shown in Table B-9.

Table B-9.
Decomposition Models

Model	Formulation of forecast		
Additive model	Actual	=	$T + S + C + I$
	Forecast	=	$T + S + C$
Multiplicative model	Actual	=	$T \bullet S \bullet C \bullet I$
	Forecast	=	$T \bullet S \bullet C$

Example

To solve for T, S, and C involves considerable manipulation of the data and special algorithms. For example, the T component could be found with the linear trend model previously described. Because this method is tedious, we will not apply it to our example; rather, we refer you to a book by Ellis and Nathan.[1]

[1] Ellis, Dennis, and Jay Nathan, *A Managerial Guide to Business Forecasting*, Graceway Publishing Co., New York, 1990.

Seasonal Smoothing Models

Description

The Winter's three-parameter linear exponential smoothing model represents an extension of the exponential smoothing models so that they can handle both trend and seasonal influences. Somewhat like the decomposition method, this approach develops three components – level, trend, and seasonal – and uses them to formulate a forecast.

Model

The following four equations for the model are taken from Ellis and Nathan's book.[2]

Level:	L_t	=	$\alpha(Y_t/S_{t-m}) + (1-\alpha)(L_t - 1 + T_{t-1})$
Trend:	T_t	=	$\beta(L_t - L_{t-1}) + (1-\beta)(T_{t-1})$
Seasonal:	S_t	=	$\gamma(Y_t/L_t) + (1-\gamma)S_{t-m}$
Forecast:	F_{t+k}	=	$(L_t + kT_t)S_{t-m}$

where Y_t = the actual observation of the data at time period t

L_t = the smoothed estimate of the level of the data at period t. L_{t-1} means the smoothed estimate of the level of the data last period (lagged 1 period). L is measured in the same units as the data.

T_t = the smoothed estimate of the trend influence on the data at period t. T_{t-1} means the smoothed estimate of the trend influence of the data last period. T is measured in the same units as the data.

S_{t-1} = the smoothed estimate of the seasonal influence on the data at period t. S_{t-m} is the smoothed estimate of the seasonal influence in periods ago; m is the length of seasonality. S is measured in index number form; it is a seasonal index.

F_{t+k} = the forecast for k periods ahead of the present. For example, with monthly data, F_{t+1} is a two month ahead forecast for period t.

---

[2] Ellis, ibid.

α, β, and γ	=	the weighting factors for level, trend, and seasonal, respectively. Best results are obtained when these weights are between 0 and 1.

Again, the tediousness of this method precludes us from applying it to our example, but we refer you to Ellis and Nathan's book which presents an example.[3]

Box-Jenkins Method

Description

The Box-Jenkins Method is not a model per se but is an approach to forecasting complex situations in which the data pattern is not evident. Ellis and Nathan list the steps in the approach as follows:[4]

- Decide which types of models to consider
- Identify which models will fit the data
- Estimate the necessary coefficients (of the models)
- Diagnose the model.

The development of the Box-Jenkins method is complex, and a thorough knowledge of its use requires higher-order mathematics beyond that of forecasting specialists.

Models

Three types of models use the Box-Jenkins method:

- *Autoregressive model (AR).* The AR model postulates that the current value of a variable is the weighted linear sum of past values plus some error term.
- *Moving average (MA).* The MA model postulates that the current value of a variable is a weighted linear relationship of past error terms and the current random term.
- *Integrated autoregressive moving average (ARMA).* The ARMA model postulates that the current value of the variable is the combination of the AR and MA models.

---

[3] Ellis, ibid.
[4] Ellis, ibid.

Makridakis and Wheelwright and Anderson have written greater detail on the Box-Jenkins Method.[5]

COMBINED FORECASTS

Description

Not only do combined forecasts have intuitive appeal, but academics and practitioners often recommend them. Data patterns often exhibit both trend and randomness or two or more other attributes. Since different models work better with different data patterns, limiting a forecast to a single model in this case may not produce a good forecast. Using an average of two or more forecasting techniques may be better than using a 'wrong' model or a single poor forecasting model.

Unless strong evidence indicates a particular forecasting model is better than other models for a given data pattern, it might be desirable to combine the output from several models. A combined forecast is less sensitive to the specific choice of models, and it uses more information about the data pattern than a single model. The potential for large errors is reduced because the forecast is not built on a single set of assumptions. Therefore, it is safer and less risky than relying on a single model. Moreover, it is cheaper to use a combination of relatively simple methods (e.g., single exponential smoothing) than to use a single model that is more complex and requires personalized data analysis and costly data fitting.

Models

Research to date suggests that equal weighting of the various forecasts is sufficient. Our approach to combined forecasts is to take an average of the cheapest or more easily understood methods. Thus, we chose the following combinations:

- *The basic model and an eight-quarter moving average* – The average of these two models combines the overall smoothing of random demand found with the moving average with the assumption that the most recent demand is the best predictor of future demand found with the basic model.

- *The basic seasonal model and an eight-quarter moving average* – The average of these two models combines the overall smoothing of random demand found with the moving average with the seasonality assumption found with the basic seasonal model.

---

[5] Makridakis, Spyros, and Steven C. Wheelwright, *Forecasting Methods and Applications*, John Wiley and Sons, New York, 1978 and Oliver Anderson, *Time-Series Analysis and Forecasting – the Box-Jenkins Approach*, Butterworth, London, 1975.

- *The basic seasonal model and single exponential smoothing* – The average of these two models combines the seasonality assumption found in the basic seasonal model with the smoothing and emphasis on more recent demand found in the exponential smoothing model.

- *The basic seasonal model, eight-quarter moving average, and regression model* – The average of these three models combines the seasonality of the basic seasonal model with the smoothing from the moving average and the trend from the regression model.

Example

The computations of the combined forecast is simply the average of the forecasts from the individual models. The results of this averaging are summarized in Table B-10.

Table B-10.
Combined Models Sample Computations

Pattern	Model and forecast				Actual value
	Basic and moving average	Seasonal and moving average	Seasonal and single smooth	Seasonal, moving average, and regression	
Level	10	10	10	10	10
Trend	36.2	28.8	29.6	35.8	50
Seasonal	20	15	14.6	18.1	10
Cyclical	20	17.5	18.1	17.4	25
Random	24.8	21.8	21.4	11.8	36

APPENDIX C

Nonlevel-Demand, Lot-Size Algorithms

In this appendix, we describe the variety of algorithms available for determining a lot size (i.e., order quantity) for a future multiple-period demand pattern that is not stable (i.e., nonlevel demand). We used the term "lot size" because the models are most associated with determining production lot sizes. However, we will be applying them to compute order quantities.

To help understand how the models compute quantities, we present an example from Silver and Peterson's text on inventory management.[1] In the example, the MIDAS International plant in Germany is solving for the optimal production quantities to fill its customer's (MIDAS Canada's) requirements. The MIDAS Canada projected requirements for the coming year are given in Table C-1.

Table C-1.
Nonlevel-Demand Example

Month	*Sequential* number	Requirements (boxes)
January	1	10
February	2	62
March	3	12
April	4	130
May	5	154
June	6	129
July	7	88
August	8	52
September	9	124
October	10	160
November	11	238
December	12	41
Total		1,200

[1] Silver, Edward L., and Rein Peterson, *Decision Systems for Inventory Management and Production Planning*, John Wiley & Sons, New York, 1985.

For purposes of our example, we will assume the following values for variables:

- Cost per box is $10
- The cost to order is $300
- The holding rate is 23 percent or $0.01917 /$/month.

Least Unit Cost

The least unit cost (LUC) algorithm calculates the order quantity as the sum of the months of demand that have the smallest per-unit-cost. To determine the correct number of months, the algorithm starts with the first month and computes its per-unit-cost. It then proceeds to the second month and computes the per-unit-cost if the order quantity were the sum of the demand for Months 1 and 2. It continues the calculation through all 12 months. To calculate per-unit-cost, it uses the demand by month (D1), the cost to order (A), the cost of a box (C), and the holding rate (H) tabulated as shown in Table C-2:

Table C-2.
LUC Algorithm

Buy months	Order quantity	Cost to order	*Holding cost*	*Unit cost*
1	D_1	A	0	A/D1
1 – 2	D_1+D_2	A	HCD_2	$(A+HCD_2)/(D_1+D_2)$
1 – 3	$D_1+D_2+D_3$	A	$HC(D_2+2D_3)$	$(A+HCD_2+2HCD_3)/(D_1+D_2+D_3)$

Applying this algorithm to our example, we have the results for the first 9 months shown in Table C-3. In our example, buying a quantity of 585 (months 1 through 7) is the selected order quantity since it has the lowest unit cost.

Table C-3.
LUC Sample Computations

Buy months	Order quantity (units)	Ordering cost	Holding cost ($)	Total cost ($)	Unit cost ($)
1	10	300.00	0.00	300.00	30.00
1 – 2	72	300.00	11.88	311.88	4.33
1 – 3	84	300.00	16.48	316.88	3.77
1 – 4	214	300.00	92.12	391.23	1.83
1 – 5	368	300.00	209.3	509.3	1.38
1 – 6	497	300.00	332.93	632.93	1.27
1 – 7	585	300.00	434.13	734.13	1.25
1 – 8	637	300.00	503.89	803.89	1.26
1 – 9	761	300.00	694.03	944.03	1.31

LEAST TOTAL COST

The least total cost (LTC) looks at the total costs for all options for buying and selects the option with the lowest total cost. This approach could be a large task even in the case of our simple 12-month example. To illustrate, we selected a few options and show their computations in Table C-4. Of the options shown, the selected option would be the one with two buys because it has the lowest total cost.

Table C-4.
LTC Sample Computations

Option	Buy months	Order quantity (units)	Ordering cost ($)	Holding cost ($)	Total cost ($)
One buy	1 – 12	1,200	300.00	1,512.63	1,812.63
Two buys	1 – 7	585	300.00	434.13	1,287.51
	8 – 12	615	300.00	253.38	
Three buys	1 – 7	585	300.00	434.13	1,427.09
	8 – 10	336	300.00	85.10	
	11 – 12	279	300.00	7.86	
Twelve buys	1,2,3,4,5,6, 7,8,9,10,11, 12	214	12(300.00)	0.00	3,600.00

Part-Period Balancing (PPB)

Part-periods are the number of periods that parts are held in the inventory. The PPB techniques starts with the demand for the first period and continues to add demand period-by-period as long as the cumulative number of part-periods is less than the critical value found by dividing the cost-to-order by the holding cost (unit cost times holding rate per month). Table C-5 illustrates the technique for our example where the critical value is 1,565.[2] With the given critical value, we stop at 1,737 and select an order quantity for the six months or 497 units.

Table C-5.
PPB Sample Computations

Month	1	2	3	4	5	6	7
Requirement	10	62	12	130	154	129	88
Part-period	0	62	24	390	616	645	528
Cumulative	0	62	86	476	1,092	1,737	2,260

To improve the performance of PPB, a look-ahead test was developed that does the following:

If $(N)\ (R_{f+1}) < R_{f+2}$, add R_{f+1} to order

If $(N)\ (R_{f+1}) > R_{f+2}$, use initial order

where

N = the periods that requirements in period f+1 are carried if added

R_{f+1} = the requirements in the period following final period in initial order.

For our example, $R_{f+1} = 88$, N $R_{f+1} = 6 \times 88 = 528$, and $R_{f+2} = 52$. Since 528>52, we would use the initial order.

Stone offers still another version incorporating the idea of a group order quantity where the demand for a period is the group's demand and the cost-to-order in the PPB technique is replaced by a group cost-to-order.[3]

[2] 300(0.01917 • 10) = 1,565.

[3] Stone, Gerald C., *Lot Sizing for the 80s*, American Production & Inventory Control Society, Readings in Production & Inventory Control and Planning, Falls Church, Virginia, 1984, pp 40 – 43.

Wagner-Whitin Heuristic

The Wagner-Whitin (W-W) technique applies dynamic programming to determine the low-cost lot size. Like the economic order quantity (EOQ), it considers the costs of ordering and holding stock. The algorithm starts by comparing the cost of one order for the first two periods against the cost of two orders, one for each period. It selects the low-cost solution. Then, it moves to three periods, deciding between one order for all three periods or the low-cost solution for two periods plus one order for the additional period. The algorithm continues in this fashion until *N* periods are completed, where *N* periods is the preset time horizon.

Using the example from above, we start with the first 2 months and compare one versus two buys:

- *If we make one buy* – covers both months, the total cost is the sum of a $300 ordering cost and a $12 holding cost (carrying month 2 demand for one month), or $312.

- *If we make two buys* – one for each month, the total cost is a $600 ordering cost and a zero holding cost, or $600.

The best option for the first 2 months is to make one buy covering both months.

Next, we consider the best options for the first 3 months:

- *If we make one buy* – covers 3 months, the total cost is the sum of a $300 ordering cost and a $16 holding cost (carrying month 2 demand for one month and month 3 demand for two months), or $316.

- *If we make two buys* – one for Month 1 and another covering Months 2 and 3, the total cost is the sum of a $600 ordering cost and a $2 holding cost (carrying month 3 demand for one month), or $602.

- *If we make two buys* – the best option for Months 1 and 2 and another for Month 3, the total cost is the sum of a $600 ordering cost and a $12 holding cost (carrying Month 2 demand for 1 month), or $612.

The best option for the first 3 months is to the first option of make one buy covering all 3 months. The algorithm continues in this manner for some fixed number of months; i.e., that is, the algorithm does not end by itself but requires the user set a horizon.

Silver-Meal Heuristic

The Silver-Meal (S-M) technique is another heuristic for determining the low-cost lot size. The algorithm seeks to determine the lot size when the total

cost per unit of time (i.e., in our case, a month) for the duration of the order quantity is minimized. It starts by computing the total cost of a 1-month buy; that is, a buy that covers the first month of demand. Then it computes the total cost per unit of a 2-month buy; that is, a buy that covers the first two months. It continues in this fashion until the cost per month stops declining. The optimal lot size is equal to the demand over the number of months for which the lowest cost per month occurred.

The S-M technique applied to our example yields the results in Table C-6. The lot size would cover the first 4 months since Month 4 has the least total relevant cost per unit time.

Table C-6.
S-M Sample Computations

Month	Demand for month	Order cost ($)	Carrying cost for month's demand ($)	Row sum ($)	Cumulative sum ($)	Cumulative sum divided by number of months ($)
1	10	300.00		300.00	300.00	300.00
2	62		11.88	311.88	311.88	155.94
3	12		4.60	316.48	316.48	105.49
4	130		74.75	391.23	391.23	97.81
5	154		118.07	509.3	509.30	101.86

APPENDIX D

Low-Demand Level Setting Example

This appendix presents a six-item example of computing response-based levels for low-demand items. The characteristics of the items are shown in Table D-1. Assuming the Poisson distribution for the expected number of units in the pipeline (that is, in the lead time), we compute the expected number of back-orders on-hand (EBO) for each item and the benefit/cost ratio for increasing levels of spares (see Table D-2).

Table D-1.
Low-Demand Level Setting (Example Item Data)

Item	1	2	3	4	5	6
Cost	$100	$100	$500	$500	$100	$100
Annual demand	1	1	1	1	2	2
Lead time (days)	360	180	360	180	360	180

Table D-2.
Level-Setting Computations (For Example)

Item	1	2	3	4	5	6
EBO for no spares	1	0.5	1	0.5	2	1
EBO for 1 spare	0.3679	0.6065	0.3679	0.1065	0.1353	0.3679
Benefit/cost	0.0063	0.0039	0.0013	0.0008	0.0086	0.0063
EBO for 2 spares	0.1036	0.0163	0.1036	0.0163	0.5413	0.1036
Benefit/cost	0.0026	0.0009	0.0005	0.0002	0.0059	0.0026
EBO for 3 spares	0.0233	0.0019	0.0233	0.0019	0.218	0.0233
Benefit/cost	0.0008	0.0001	0.0002	0.0000	0.0032	0.0008

When we start with no spares for the six items, we have a total of six EBO's across all the items. If we wanted to cut that number in half, we start by looking for the highest benefit/cost ratio for adding one spare. Item 5 has the highest, so we would set a level of one for item 5. The new EBO total is 5.1353 - above our target of 3. So we look for the next "best" spare. Item 6 has the second highest ratio, so we add one to its level. The EBO total is now 4.5032, so we look for the

next best spare. Item 1 is next but the EBO total is 3.8711 after adding one for it. Since the next best ratio is for spare 2 for Item 5, we make Item 5's level two spares, and we continue the same way.

We end with the results shown in Table D-3 with a EBO total of 2.8310, which is less than our goal of 3, and $600 in inventory.

Table D-3.
Summary of Example Results

Item	1	2	3	4	5	6
Level	1	1	0	0	3	1

Using this procedure across all of our low-demand items, we could construct a dollar and backorder curve to show how our backorders decline as we invent money in inventory. We could use this curve to help decide what our investment should be.

APPENDIX E

Glossary

ACWT	=	average customer wait time
ALT	=	administrative lead time
AR	=	autoregressive forecasting model
ARMA	=	integrated autoregressive moving average forecasting model
BAS	=	basic forecasting model
BOA	=	basic ordering agreement
BPA	=	blanket purchase agreement
DLA	=	Defense Logistics Agency
EBE	=	expected backorders established
EBO	=	expected backorders
EOQ	=	economic order quantity
ERM	=	economic retention model
ERQ	=	economic repair quantity
IM	=	item manager
JIT	=	just in time
L4L	=	lot-for-lot (order quantity algorithm)
LFL	=	lot-for-lot (order quantity algorithm)
LHS	=	left hand side (of equation)
LORA	=	level-of-repair analysis
LTC	=	least total cost (order quantity algorithm)

LUC	=	least unit cost (order size algorithm)
MA	=	moving average forecasting model
MA4Q	=	4-quarter moving average forecasting model
MA8Q	=	8-quarter moving average forecasting model
MAD	=	mean absolute deviation
MADD	=	mean absolute deviation for reparable demand not satisfied by regeneration
MADG	=	mean absolute deviation for reparable regenerations
MADLT	=	mean absolute deviation over lead time
MAPE	=	mean absolute percentage error
ME	=	mean error
MPE	=	mean percentage error
MSE	=	mean square error
MSRT	=	mean system response time
PLT	=	procurement lead time
PLTQ	=	procurement (acquisition) lead-time quantity
PPB	=	part-period balancing (order quantity algorithm)
QPS	=	quantity per system
RBS	=	readiness-based sparing
RCT	=	repair cycle time
RCTQ	=	repair cycle time quantity
REGR	=	linear regression forecasting model
RHS	=	right hand side (of equation)
SBAS	=	basic seasonal forecasting model
SCB	=	Supply Center Baltimore

SCCB	=	Supply Center Curtis Bay
SD	=	standard deviation
SES1	=	single-exponential smoothing with alpha = 0.1 forecasting model
SES2	=	single-exponential smoothing with alpha = 0.2 forecasting model
SLQ	=	safety level quantity
TVC	=	total variable cost
USAGE	=	simple usage forecasting model

Appendix F

Bibliography

Abraham, Bovas, and Johannes Ledolter, *Statistical Methods for Forecasting,* John Wiley & Sons, New York, 1983.

Anderson, Oliver, *Time Series Analysis and Forecasting – The Box-Jenkins Approach,* Butterworth, London, 1975.

Cameron, William E., "Case Studies of Supply Inventory Models," Air Force Institute of Technology Thesis, September 1988.

Cost to Order and Cost to Hold Study, USCG Supply Center Curtis Bay, Fleet Support Operations Branch, 1 June 1992.

DLA Economic Retention/Returns Limits Study. Operations Research and Economic Analysis Office, Defense Logistics Agency, September, 1986.

Ellis, Dennis, and Jay Nathan. *A Managerial Guide to Business Forecasting,* Graceway Publishing Co., New York, 1990.

Guidance for Implementation of Time-Weighted Essentiality-Weighted Requisitions Short Variable Safety Level, Operations Research and Economic Analysis Office, Defense Logistics Agency, July 1972.

Hadley, G., and T.M. Whitin, *Analysis of Inventory Systems,* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1963.

Kiebler, Kelvin, Larry Klapper, and Don Frank, *Army Depot Maintenance: More Effective Use of Organic and Contractor Resources,* Logistics Management Institute, Report AR803R1, June 1990.

LaForge, R. Lawrence, "MRP and the Part-Period Algorithm," *Journal of Purchasing and Materials Management,* Winter 1982, pp 21 – 26.

Makridakis, Spyros and Steven C. Wheelwright, *Forecasting Methods and Applications,* John Wiley & Sons, New York, 1978.

Multiple Forecasting Techniques, Operations Research and Economic Analysis Office, Defense Logistics Agency, December 1990.

O'Malley, T.J., *The Aircraft Availability Model: Conceptual Framework and Mathematics,* Logistics Management Institute, Report AF201, June 1983.

Plossl, George, W. *Production and Inventory Control: Principles and Techniques*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1985.

Presutti, Victor J., Jr., and Richard C. Trepp, "More Ado About Economic Order Quantities (EOQ)," *Naval Research Logistics Quarterly*, Vol. 17, No. 2, June 1970, pp 243 – 251.

Silver, Edward L., and Rein Peterson, *Decision Systems for Inventory Management and Production Planning*. John Wiley & Sons, New York, 1985.

Sims, Robert L., and Dennis L. Zimmerman, *Multiple Cost EOQ Study*, Operations Research and Economic Analysis Office, Defense Logistics Agency, December 1989.

Smith, Bernard T., *Focus Forecasting: Computer Techniques for Inventory Control*. CBI, Boston. 1978.

Stockage Policy Analysis, Office of the Assistant Secretary of Defense (Manpower, Reserve Affairs, and Logistics), Final Report, Annex A, Part 3, 31 August 1980, pp 1 – 12 to 1 – 14.

Stone, Gerald C., *Lot Sizing for the 80s*, American Production & Inventory Control Society, Readings in Production & Inventory Control and Planning, Falls Church, Virginia, 1984, pp 40 – 43.

Zimmerman, Dennis, Robert L. Sims, and Andrew A. Pelletz, *Cost-to-Hold Methodology*, Operations Research and Economic Analysis Office, Defense Logistics Agency, February 1993.

REPORT DOCUMENTATION PAGE

Form Approved
OPM No.0704-0188

1. AGENCY USE ONLY ***(Leave Blank)***

2. REPORT DATE
October 1993

3. REPORT TYPE AND DATES COVERED
Final

4. TITLE AND SUBTITLE
Improved Inventory Models for the United States Coast Guard Requirements Determination Process

5. FUNDING NUMBERS
C MDA903-90-C-0006
PE 0902198D

6. AUTHOR(S)
George L. Slyman
Dennis L. Zimmerman

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Logistics Management Institute
6400 Goldsboro Road
Bethesda, MD 20817-5886

8. PERFORMING ORGANIZATION REPORT NUMBER
LMI– CG201RD6

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Office of Engineering, Logistics and Development
Headquarters, U.S. Coast Guard
2100 Second Street, S.W.
Washington, DC 20593

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
A: Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 words)***

Requirements determination is the process by which the Coast Guard supply system forecasts future customer demands and sets levels of inventory to satisfy those demands. Currently, the Coast Guard is in the process of modernizing its inventory-management system at its nonaviation supply centers. To support them, the Logistics Management Institute undertook a study to update the inventory models in the Coast Guard's requirements determination process.

Although different types of inventory-management models deal with different decisions, they are all interconnected. For example, the demand forecast from a forecasting model feeds into the computation of stockage levels performed by other models. It is not enough to improve demand forecasts if the levels computations that use those forecasts are not updated to capitalize on the improvements.

This report documents the LMI study of how the Coast Guard can update all of its inventory models. We make specific recommendations on which inventory models we believe the Coast Guard should use in its inventory-management system and we present those recommendations at the end of each chapter.

14. SUBJECT TERMS
USCG; Requirements Determination Process, Inventory Models, Inventory Management, Materiel Management

15. NUMBER OF PAGES
138

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298, (Rev. 2-89)
Prescribed by ANSI Std. 239-18
299-01